

The Foundation provides awards for research in the sciences and engineering. The awardee is wholly responsible for the conduct of such research and preparation of the results for publication. The Foundation, therefore, does not assume responsibility for such findings or their interpretation.

The Foundation welcomes proposals on behalf of all qualified scientists and engineers, and strongly encourages women, minorities, and persons with disabilities to compete fully in any of the research-related programs described here. In accordance with federal statutes, regulations and NSF policies, no person on grounds of race, color, age, sex, national origin, or disability shall be excluded from participation in, denied the benefits of, or be subject to discrimination under any program or activity receiving financial assistance from the National Science Foundation.

Facilitation Awards for Scientists and Engineers with Disabilities provide funding for special assistance or equipment to enable persons with disabilities (investigators and other staff, including student research assistants) to work on an NSF project. See the program announcement, or contact the program coordinator at (703) 306-1636.

The National Science Foundation has TDD (Telephonic Device for the Deaf) capability, which enables individuals with hearing impairment to communicate with the Foundation about NSF programs, employment, or general information. To access NSF TDD dial (703) 306-0090; for FIRS, 1-800-877-8339.

Catalog of Federal Domestic Assistance Number 47-070; Computer and Information Science and Engineering.

Preface

The Computer and Information Science and Engineering (CISE) Directorate, under the direction of an Assistant Director, consists of the following six divisions and offices: Advanced Scientific Computing (ASC) Division, Computer and Computation Research (CCR) Division, Cross-Disciplinary Activities (CDA) Office, Information, Robotics and Intelligent Systems (IRIS) Division, Microelectronic Information Processing Systems (MIPS) Division, and the Networking and Communications Research and Infrastructure (NCRI) Division.

The **Microelectronic Information Processing Systems Division (MIPS)** supports research on novel computing and information processing systems including signal processing. Emphasis is on experimental research, technology-related research and particularly the critical link between conceptualization and realization for integrated systems. Technologies include VLSI, ULSI, OPTICAL, OPTO-ELECTRONIC, INTER-CONNECTION and other emerging technologies. The focus is on research pertaining to hardware systems and their supporting software, including: experimental research involving these new systems; infrastructures, environments, tools, methodologies and services for rapid systems prototyping; design methodologies and tools; technology-driven and application-driven systems architectures; and fabrication and testing of systems. For signal-processing systems, research on algorithms and architectures relating to these new technologies that have promise for real-time computing is emphasized.

The purpose of this Summary of Awards for the MIPS Division is to provide the scientific and engineering communities with a summary of those grants funded in Fiscal Year 1995. This report lists only those projects funded using Fiscal Year 1995 dollars and hence may not list some multi-year awards initiated prior to Fiscal Year 1995.

Similar areas of research are grouped together for reader convenience. The reader is cautioned, however, not to assume that these categories represent the totality of interests of each program, or the total scope of each grant. Projects may bridge several programs or deal with topics not explicitly mentioned herein. Thus, these categories have been assigned administratively and for the purpose of this report only.

In this document, grantee institutions and principal investigators are identified first. Award identification numbers, award amounts, and award durations are enumerated after the individual project titles. Within each category, the awards are listed alphabetically by state and institution.

Readers wishing further information on any particular project described in this report are advised to contact the principal investigators directly.

Bernard Chern
Division Director
Microelectronic Information
Processing Systems Division

Table of Contents

Preface	iii
Table of Contents	v
The MIPS Division	vii
MIPS Directions	ix
MIPS Staff	xi
Summary	xiii
 Design, Tools and Test	 1
The Program	1
Initiatives and Opportunities	2
Awards	3
Design Automation	3
Manufacturing Test	11
Simulation	14
Other	15
 Microelectronic Systems Architecture	 19
Program Description	19
Initiatives and Opportunities	20
Awards	21
Technology Driven Architecture	21
Application Driven Architecture	30
Workshops and Conferences	34
 Circuits and Signal Processing	 35
The Program	35
Awards	37
Analog (Mixed Analog/Digital) Signal Processing	37
One-Dimensional Digital Signal Processing	39
Image and Multidimensional Digital Signal Processing	42
Statistical Signal and Array Processing	48
Other	52

Experimental Systems	55
The Program	55
Initiatives and Opportunities	56
Awards	57
Storage Hierarchies and Input/Output Systems	57
General Purpose Computing	58
Application Specific Computing	61
Other	65
Systems Prototyping and Fabrication	67
The Program	67
Initiatives and Opportunities	68
Awards	69
Prototyping	69
Microelectronic	69
MEMS, SFF, MCM	72
Education	73
Infrastructure	74
Index of PYI, NYI, CAREER and PFF Investigators	75
Index of Principal Investigators	77
Index of Institutions	81

Microelectronic Information Processing Systems

The MIPS Division

The area of Computing Systems, which involves the structure of computers, is central to MIPS today and will be even more so in the future. This is a core area of computer science and engineering and in the 1990's encompasses much more than just hardware. Computing systems deals with computer architecture, hardware implementation, system software (operating systems and compilers), networking, and data storage systems. The advent of gigabit networks, high performance microprocessors and parallel systems is dramatically impacting research on systems level architecture of high performance computing systems.

The emphasis in MIPS is on REAL SYSTEMS i.e. physically realizable. Special weight is placed on design, prototyping, evaluation, and novel use of computing systems and on the tools needed to design and build them. This involves technology driven and application related research, experimental research and theoretical studies. The MIPS programs support research on: high level design (design automation and CAD tools); systems level architecture studies; experimental systems research projects which build and evaluate HARDWARE/SOFTWARE SYSTEMS; signal processing algorithms and systems; knowledge of applications; methodologies, tools and packaging technologies for rapid prototyping at the system level; and infrastructure needed to support MIPS' educational and research activities, e.g. MOSIS.

The Programs

Design, Tools and Test Program

The objective of the Design, Tools and Test Program is to obtain fundamental knowledge about the complete design cycle for integrated circuits and systems from conception through manufacturing and operational test. Emphasis is on integrating all aspects of the cycle, and automating the design and testing processes. There are four topical research areas within the Program. These are: Theoretical Foundations, Design Automation and Tools, Manufacturing Test, and Design Simulation.

Systems Prototyping and Fabrication Program

Supports research on technologies, tools, and methodologies needed for the prototyping of experimental information processing systems and for Microelectronics Education. Issues that arise in rapid system prototyping are explored, including use of new packaging techniques such as multichip modules, and such systems issues as interfacing and standards. Support is also provided for new prototyping services. Basic research necessary to model, simulate, measure, automate and improve the microfabrication process is supported. Microelectronics Education support includes workshops, conferences, development of curriculum and courseware materials, and educational support services such as those for FPGA's and fabrication (MOSIS).

Microelectronic Systems Architecture Program

Supports basic research on computing systems and methods for their design. Computing Systems deals with computer architecture, hardware implementation, systems software, networking, and data storage. Research is encouraged on the fundamental aspects of computing systems architectures and scientific design methods that better utilize existing or emerging technologies, support systems software or address important applications whose computational requirements cannot be met by conventional architectures. The program emphasizes physically realizable systems and, when necessary, limited proof-of-concept prototyping.

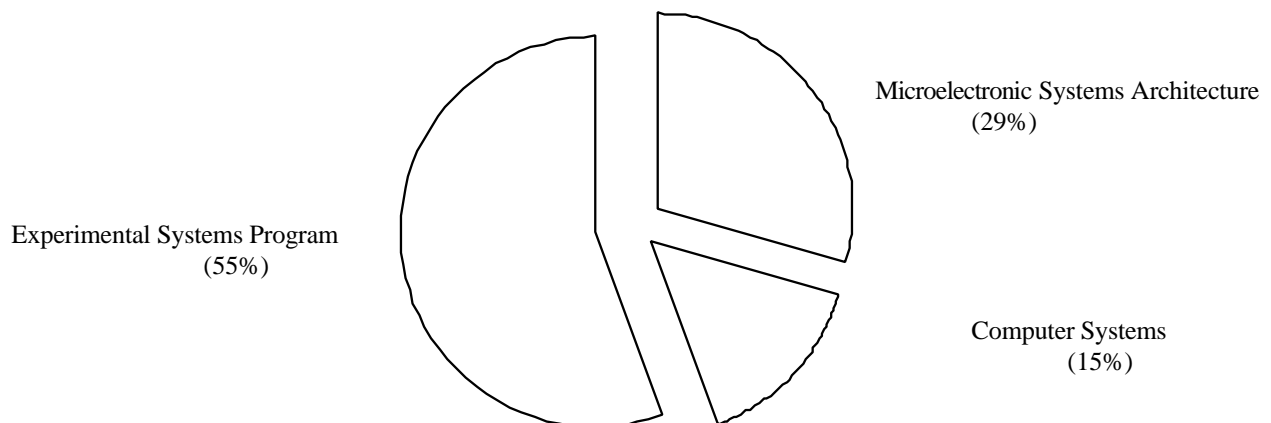
Circuits and Signal Processing Program

Supports research on circuit theory and analog and digital signal processing. The emphasis is on modern signal processing, stressing the impact of VLSI, including areas such as: signal representation, filtering, novel algorithms, special-purpose hardware, and real-time computing. Circuit theory research encompasses such activities as nonlinear, discrete-time, analog and hybrid circuits, and analog/digital conversion.

Experimental Systems Program

Supports research experiments that involve building and evaluating information processing and computing systems. These are goal-oriented projects usually undertaken by teams of designers, builders, and users. The building of a system must itself represent a major intellectual effort, and offer advances in our understanding of information systems architecture by addressing significant and timely research questions. The system prototypes being built should be suitable for exploring applications and performance issues.

Basic research in Computer Architecture and Computing Systems is supported by the National Science Foundation primarily through three programs in the CISE Directorate: the Microelectronic Systems Architecture Program and the Experimental Systems Program in the MIPS Division; and the Computer Systems Program in the CCR Division. The following pie chart shows the relative support through these three programs.



MIPS Directions

MIPS' planning takes into account advances in technology and new knowledge, and the need for closer ties of computer science and engineering to real world applications. Greater emphasis is now placed on **complete systems**: with a broad and coherent research program in new systems architectures, automated design, and design tools to aid in research and development of high performance architectures.

The MIPS role in the High Performance Computing and Communications (HPCC) Program focuses on the support of:

- Basic research (hardware & system software) on new high performance computer architectures and computing systems and on "Computer Science Challenges";
- Development of tools & CAD frameworks for their design, analysis and realization;
- Algorithm development and computational techniques for "grand and national challenge" problems in the areas of research supported by MIPS.

Research on high performance computing systems is responsive to such major drivers as: technology, applications and new ideas. To make advances in this area that can be effectively exploited, requires that experimental systems be built quickly and cheaply and new kinds of design tools be developed and supplied to the research community to enable this to occur. These prototype systems can be used to evaluate new computing architectures by subjecting them to real applications which provide believable tests of novel ideas and performance. Only by constructing prototypes, performing measurements and evaluating performance, can we realistically gauge the interaction between a new computing system, its applications, and its users.

We see the need to:

- Work more closely with the applications as we move toward higher performance computing to understand the computing needs of these applications.
- Integrate advanced packaging technology into computing system design and explore the system level tradeoffs arising in the design of high performance computing systems.
- Develop the necessary infrastructure and human resources in the computing systems area, especially the education of students able to design and build hardware/software systems.
- Develop new services, tools, and methodologies for universities to utilize new fabrication and device technologies in order to do rapid system prototyping for electrical and mechanical systems essential for experimental research on these increasingly complex systems.
- Support geographically distributed collaborative novel computing system design requiring expertise from many areas (e.g., architecture, software, storage technology, I/O, applications, etc.).
- Develop a new generation of systems level design tools which have increased functionality, are highly automated, and accept high level specifications as inputs.

MIPS Directions (continued)

Rapid Prototyping & Design

There is a need to be able to design and build "prototype" systems quickly, and easily in both universities and industry in order to test the validity of a concept and to shorten the design-build-evaluate cycle. With the advent of High Performance computing systems and high speed networks, this presents the opportunity to do prototyping using distributed design and manufacturing.

Rapid prototyping allows designs created with CAD systems and expressed in an all-digital format to be "prototyped" either by simulating their performance and/or by constructing limited numbers of actual parts to ascertain performance.

We feel this is an important research area involving 3 dimensional electro-mechanical objects, and so MIPS was one of the major participants in a Program Announcement on Rapid Prototyping: Virtual and Physical issued by the Foundation in April 1994.

In May, 1994 the MIPS Division supported a Workshop to examine design and prototyping as done in VLSI and in the mechanical world. The success of the VLSI revolution, leads one naturally to ask the question **what can we learn from that success, and how might it be applicable to the design of electro-mechanical systems?** This was the central issue the workshop examined. A number of Findings and Recommendations were put forward by the workshop¹. In particular, it recommended the active investigation of Solid Freeform Fabrication (SFF) and Micro electro-mechanical Systems (MEMS) mechanical implementation technologies where the VLSI experience will be most relevant and provides a good starting point to apply VLSI-like Systems Design Methodologies.

The MIPS Division then supported a workshop at Carnegie Mellon University on June 5-6, 1995 dealing with Design Methodologies for Solid Freeform Fabrication. This workshop was convened to examine in more detail existing design methodologies being used in SFF and to determine whether substantial benefit would arise from research into applying key elements from the VLSI experience to SFF rapid Prototyping technologies. The following SFF Technologies were examined: Stereolithography, Laser-based SFF, Shape Deposition Manufacturing, 3D Printing, and Lamination. The Executive Summary of the resulting report contains findings and recommendations addressing SFF Taxonomy, Design Hierarchy and Design Languages.²

On November 12-15, 1995, a workshop supported by NSF was held at the California Institute of Technology to discuss and explore the research issues involved in developing Structured Design Methodologies for Micro-electro-mechanical Systems (MEMS) which would emphasize the clean separation between design and fabrication. The report of this workshop is expected to be ready early in 1996.

As a result of these workshops, which provided an opportunity to examine more closely today's SFF and MEMS design methodologies, the MIPS Division sees the need to support research on extending those methodologies to incorporate a "clean" separation between design and fabrication. This looks to be achievable (as was done in VLSI), by adopting a generic model of the fabrication process, incorporating increased use of digital interface descriptions between design and fabrication, design rules, levels of abstraction and CAD tools, and improved languages for the description of 3-D objects.

¹ New Paradigms for Manufacturing (NSF 94-123)

² Report is in printing. Draft report is available from CMU Engineering Design Research Center (attention Dr. Susan Finger).

Division of
Microelectronics Information Processing Systems

MIPS Staff

Bernard Chern
Division Director
bchern@nsf.gov

John R. Lehmann
Deputy Division Director
jlehmann@nsf.gov

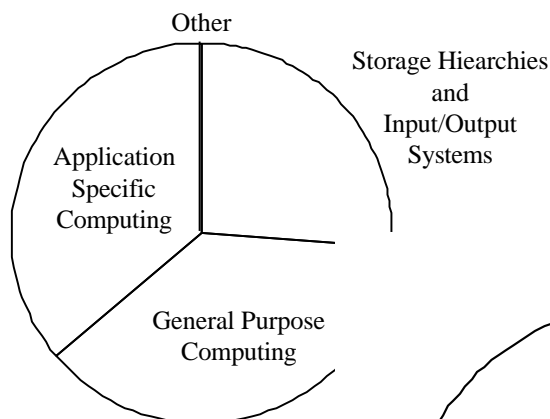
Program	Program Director	Net Address
Design, Tools and Test	Robert B. Grafton	rgrafton@nsf.gov
Microelectronic Systems Architecture	Lionel M. Ni	lni@nsf.gov
Circuits and Signal Processing	John H. Cozzens	jcozzens@nsf.gov
Experimental Systems	Michael J. Foster	mfoster@nsf.gov
Systems Prototyping and Fabrication	John Staudhammer	jstaudha@nsf.gov

The address and telephone number for all of the above:

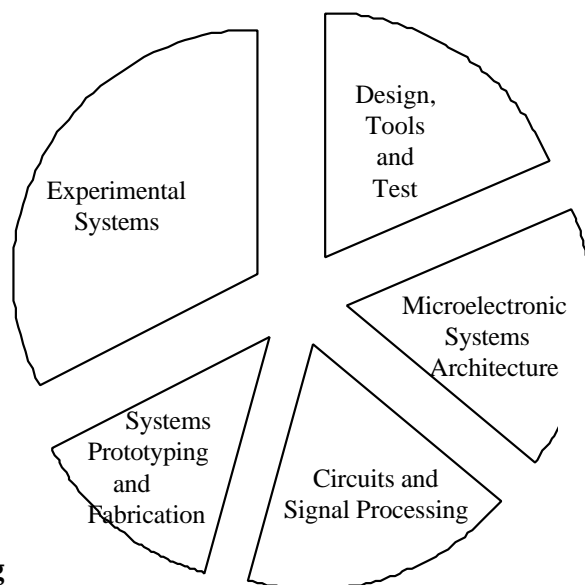
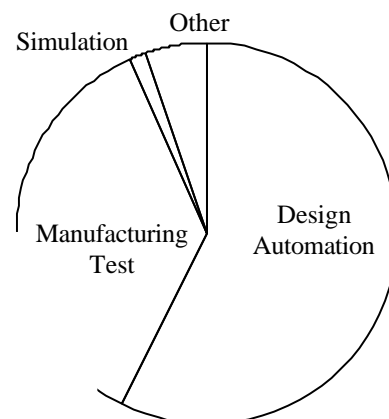
National Science Foundation
4201 Wilson Blvd.
Arlington, Virginia 22230
(703) 306-1936

MICROELECTRONIC INFORMATION PROCESSING SYSTEMS DIVISION (FY 1995)

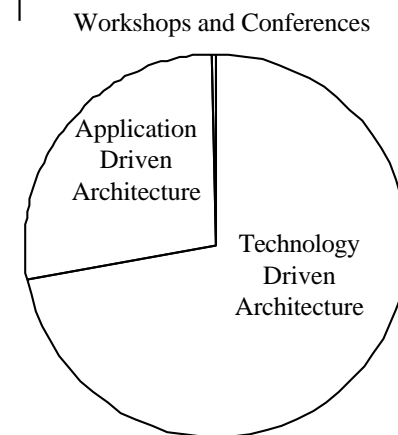
Experimental Systems



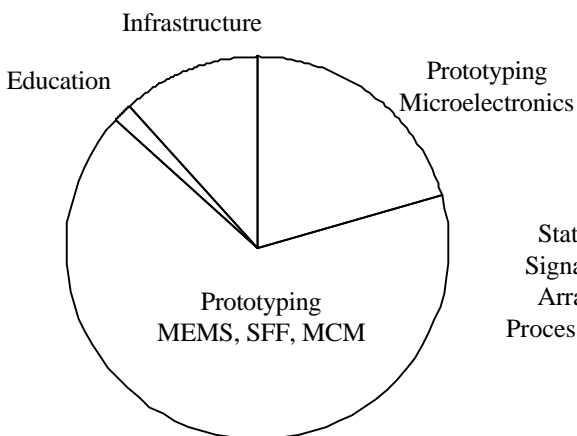
Design, Tools and Test



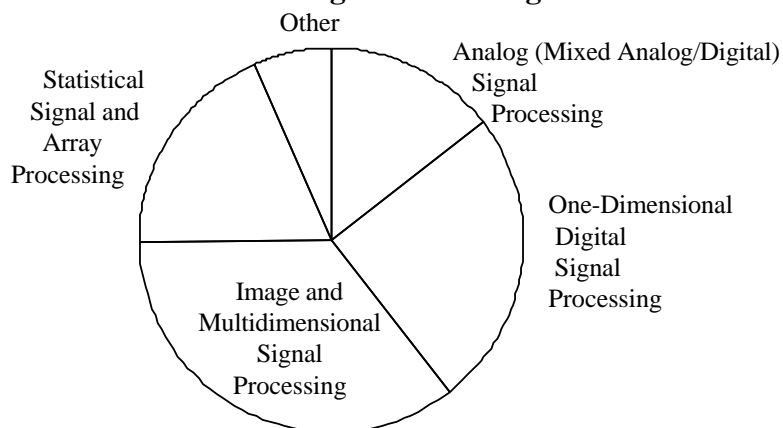
Microelectronic Systems Architecture



Systems Prototyping and Fabrication



Circuits and Signal Processing



Summary

	Number of Awards	Dollars
Design, Tools and Test	66	\$4,693,973
Design Automation	40	\$2,729,973
Manufacturing Test	16	\$1,647,556
Simulation	1	\$62,500
Other	9	\$253,776
Microelectronic Systems Architecture	51	\$4,085,077
Technology Driven Architecture	35	\$2,939,613
Application Driven Architecture	15	\$1,130,464
Workshops and Conferences	1	\$15,000
Circuits and Signal Processing	59	\$4,603,137
Analog (Mixed Analog/Digital) Signal Processing	8	\$666,102
One-Dimensional Digital Signal Processing	11	\$1,123,468
Image and Multidimensional Digital Signal Processing	19	\$1,565,891
Stistical Signal and Array Processing	12	\$917,309
Other	9	\$330,367
Experimental Systems	34	\$8,306,180
Storage Hiearchies and Input/Output Systems	7	\$2,173,211
General Purpose Computing	12	\$3,068,618
Application Specific Computing	14	\$3,059,351
Other	1	\$5,000
Systems Prototyping and Fabrication	23	\$3,035,062
Prototyping	19	\$2,633,889
Microelectronic	11	\$617,093
MEMS, SFF, MCM	8	\$2,016,796
Education	3	\$51,123
Infrastructure	1	\$350,000

This summary data includes funds designated for special Foundation initiatives, and equipment matching funds from the Office of Cross-Disciplinary Activities. It does not include program funds use to support Intergovernmental Personnel Act employees, their travel costs, or costs of travel of review panelists and site visitors.

Design, Tools and Test

Dr. Robert B. Grafton, Program Director
(703) 306-1936 rgrafton@note.nsf.gov

The Program

The Design, Tools and Test Program supports basic research in Electronic Design Automation (EDA). Findings of the research contribute to fundamental knowledge about the complete EDA design cycle for integrated circuits (ICs) and systems, from conception through manufacturing test.

Emphasis is on finding methodologies needed to design, verify, simulate, and test the more complex and heterogeneous systems of the future. Discovering these new methodologies, with their appropriate abstractions to simplify and expedite the process, will provide the basis for the development of the next generation of EDA tools. Additional focus is on system level design issues and the myriad tradeoffs that must be made in developing a balanced, manufacturable, and profitable system design.

Research in EDA tools needs to address derivation of appropriate models for the design domains, technologies, and physical effects of interest and concern today. This research, while fundamental in nature, must have roots in the real world of electronic design. Thus, researchers often have ties to industry so as to keep abreast of changing industry concerns.

Topics in EDA tools research can be positioned in a three-dimensional space where the three axes represent the traditional EDA area, the application domain, and the enabling implementation technology. This framework provides a way to assess the potential impact and difficulty of the problems addressed.

DRIVING FORCES

The technology of VLSI/ULSI circuits changes rapidly. It is now possible to put up to 5 million transistors on a chip and operate at speeds in the hundreds of Mhz range. The demand for computing systems of greater complexity and higher performance is high. Interfaces, applications, and advanced computing systems require ever more sophisticated, high quality, and trusted VLSI electronics. Thus, there is substantial need for a new generation of EDA tools for IC and system design and test.

The manufacturing process is becoming so flexible that designers can specify and control device parameters. This can be a new approach for achieving low power, high speed, reliable designs. To do this, the EDA approach, dealing with design abstractions, and the Technology CAD (TCAD) research, dealing with devices and materials, need to be integrated. Design methodologies that draw from both to these areas are sought.

Competitiveness requires that the product design cycle must be reduced, and new mechanisms for rapid verification and prototyping of designs must be developed. This, in turn, requires sophisticated EDA tools which enable design engineers to find viable solutions to design problems in a highly complex design space. Additionally, designers must account for factors, such as cost and operating environment.

The changing technology, manufacturing cost and competitiveness dictate the major research problems of finding methodologies to: Rapidly design a complex IC system; Verify the resulting design; and Test the manufactured product.

GENERAL TOPICS

High operating and switching speeds make electrical effects on the chip an important factor, thus causing a new dimension in EDA tool research. It is necessary to use an integration of analytical and experimental methodologies to achieve suitable designs. Research is supported on mixed signal and digital technologies within the areas of theory, design automation, tools, manufacturing test, and simulation.

Theoretical research explores computational models and algorithms for design and testing methods in advanced technologies and in mixed signal and asynchronous systems.

Design automation investigates algorithms, tools and analysis methods for design at several levels, including physical and high level synthesis. Goals are performance, testability, synthesis, and verifiability of designs. Integration of software with hardware aspects are part of system synthesis.

Manufacturing test research includes work on evaluation of new manufacturing test methods; algorithms for test pattern generation for test of large IC designs; models for detection of realistic failure modes; and integrating testing with other design activities.

Numerical and symbolic simulation are necessary parts of the design evaluation and analysis. Of special interest is the application of computer science, numerical analysis and electrical engineering knowledge to computation of electrical effects in designs with small feature size, high speed, and mixed signal capabilities.

Initiatives and Opportunities

Some special opportunities available through the DTT Program include:

RESEARCH EXPERIENCES FOR UNDERGRADUATES (REU)

This is an opportunity to add one or two undergraduates to a grant to take part in the research activity. Since the goal is to interest promising undergraduates in a scientific research career, their participation should be meaningful to the student as well as helping the research progress.

Grant supplements are generally for one or two students, with preference given to members of groups under-represented in engineering.

SMALL GRANTS FOR EXPLORATORY RESEARCH (SGER)

This provides for a one year grant to explore a research topic that, for some reason, may be inappropriate for a normal grant submission.

SOFTWARE CAPITALIZATION GRANTS

Software capitalization grants or supplements are to facilitate development of well-documented, useable prototypes of EDA research-developed tools, and their distribution to designers and other researchers.

Awards

Design Automation

Stanford University; Giovanni De Micheli; *Logic Synthesis of Low-Power Circuits*; (MIP-9421129); \$66,365; 12 months.

This research is on techniques and tools for automated logic design of low-power, semi-custom circuits. Digital circuits are specified as models in hardware description languages that can be readily compiled into finite-state machines. The latter are described by transition diagrams or by synchronous logic networks. These models are then used to solve logic synthesis problems in encoding sequential circuits, restructuring logic networks, and library binding. Tools for a comprehensive EDA system for low-power design are being developed.

University of California-Los Angeles; Jason Cong; *NYI: Synthesis and Mapping in Lookup-Table Based FPGA Designs*; (MIP-9357582 A002); \$62,500; 12 months.

Synthesis and Mapping are necessary steps in designing Field Programmable Gate Arrays (FPGAs), and this research is a systematic study of these problems in look-up table based FPGA designs. Questions being investigated include: effects of node duplication, effects of depth relaxation, and effects of logic resynthesis during mapping and system-level partitioning mapping in multiple chip FPGA designs. Algorithms being developed are:

1. for computing optimal or near optimal synthesis or mapping solutions under a given objective function; and
2. for synthesizing a set of mapping solutions of smooth trade off between various design objectives, such as area, delay and routability.

Computational methods for integrating the synthesis and mapping steps in FPGA designs based on combinatorial and Boolean optimization techniques are being investigated.

University of California-Los Angeles; Andrew B Kahng; *NYI: Synthesis of High-Speed, High-Complexity VLSI Systems*; (MIP-9257982 A003); \$62,500; 12 months.

The unifying theme of this research is that the underlying geometries, embedding dimensions and topological representations of CAD designs, together afford a perspective for effective algorithm design. The research is in three areas:

1. Performance-driven synthesis at various levels of design, including clustering for problem decomposition and fast hierarchical placement,

and estimation of intrinsic resource requirements via topological criteria.

2. Assessment of design problem complexity based on the interaction between topology of neighborhood structures and scaling geometry in the associated cost surfaces. This includes time-bounded stochastic optimizations, and a non-monotone class of annealing methods.
3. Capturing the underlying physics of high-speed devices and interconnects while maintaining algorithmically tractable formulations. Examples are a unified routing tree optimization and the separation of the interconnect topology from subsequent geometric embedding.

University of California-San Diego; Chung-Kuan Cheng; *Circuit Partitioning*; (MIP-9315794 A002); \$24,923; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$49,846).

This project is an investigation into a new generation of hierarchical partitioning methods motivated by the need to;

1. cope with high circuit complexity,
2. improve a system's performance under I/O pin count constraints, and
3. control intermodule delay in order to optimize system performance.

Previous partitioning research is being extended by adopting different circuit models including petri nets, data flows, and state machines, by improving the efficiency and effectiveness of the methods, and by deriving theoretical results on variations of partitioning formulations. Partitioning methods are being applied to netlist mappings for various hardware emulation machines and to find potential applications of these methods to VLSI design problems.

University of California-Santa Barbara; Forrest Brewer; *Production Language Based High-Level Synthesis*; (MIP-9320752 A001); \$52,249; 12 months.

The overall goal of this research is to create a new class of synthesis tools which address the design of complex controller-data path machines under constraints of pre-defined interfaces. Because the results will be integrated with commercial EDA design tools, a design output format, which can be simulated and allows automated re-design of selected portions of the design, is being developed. In a second task, approximate sequential reachability analysis is being used to design and implement algorithms for optimizing and partitioning controller designs. The third task is to explore scheduling algorithms for both the control and data-path portions of the design. Finally, an optimizing compiler is being built. It contains algorithms which solve encoding issues for high performance designs, and performs re-scheduling of the data-path operations to minimize

required resources while maintaining design behavior.

University of California-Santa Barbara; Malgorzata Marek-Sadowska; *Layout and Logic Design*; (MIP-9419119); \$100,000; 12 months; (Joint support with the Systems Prototyping and Fabrications Program - Total Grant \$140,000).

This research is on layout driven synthesis, i. e. the intersection of logic synthesis and physical design. The focus is on restructuring logic networks in synthesized digital systems. Four topics, which meet the goals of improving routing efficiency or power consumption, are being investigated. These are:

1. Incremental logic resynthesis to control wiring,
2. Coupling wiring with logic restructuring and finding optimizations to eliminate wiring overflows.
3. Use of generalized Reed-Muller forms to analyze logic as an aid to: - designing cell libraries and for technology mapping, - developing new multi level optimization techniques, - designing networks of provably good testability.
4. Develop new methods for power optimization, at the technology independent and technology dependent levels in logic synthesis, and also find better routing tools to handle power constraints.

University of Southern California; Peter A Beerel; *CAREER: Computer-Aided Design Tools for Asynchronous Circuits*; (MIP-9502386); \$125,589; 36 months.

This research is on the development of methods for synthesis and verification of asynchronous circuits. The approach is to extend approximation techniques for state exploration, that have been successfully applied to speed-independent circuits, to semi-custom timed circuits. Tools and techniques being developed include:

1. estimators of circuit area, performance, and power at architectural, gate, and circuit levels;
2. a synthesis tool that accepts behavioral specifications and derives gate-level circuits; and
3. a hierarchical verification tool for timed circuits. The verification and synthesis tools are integrated so as to support advanced logic optimizations.

University of Southern California; Massoud Pedram; *NYI: Low Power VLSI Design*; (MIP-9457392 A001); \$50,000; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$100,000).

This research investigates modeling and estimation

of power consumption as well as techniques for minimizing power at the various levels of design abstraction (layout, logic, register- transfer and behavioral levels). Principles and methods to guide the design of power efficient electronic systems are being explored; and the impact of availability of low-power design techniques on chip, module, and system level designs is being assessed. Topics being investigated include: spatio-temporal power estimation; state assignment for low power; power dissipation in boolean networks; common subexpression extraction; and FPGA synthesis for low power.

University of Colorado; Gary Hachtel, Fabio Somenzi, Michael Lightner; *An Integrated VHDL-based Synthesis and Verification System for VLSI Systems*; (MIP-9422268); \$273,215; 12 months.

This research is on synthesis and verification of digital systems. The computational basis for the work is the binary decision diagram (BDD) data structure and extensions. In low power circuit design, new ideas in BDD technology are being used for synthesis algorithms, and to estimate power consumption via probabilistic behavior of circuits. Decomposition concepts, such as tearing, to assess properties of very large circuits are being investigated. In verification, approximate exploration ideas are being examined for use in checking equivalence of very large circuits. Hierarchical verification capabilities, where parts of the circuit are modeled at the bit level and other parts at the word level, are being examined. To provide a sound connection between high level synthesis and high level verification, as well as to validate high level VHDL descriptions, refined comparisons of non-deterministic systems, such as bi-simulation equivalence and testing equivalence and pre-orders are being explored.

Northwestern University; Majid Sarrafzadeh; *Algorithm Design for VLSI Layout*; (MIP-9207267 A003, A004); \$60,000; 12 months.

The research is in four areas of geometric algorithms for design tools. The research topics are:

1. floor planning by graph dualization;
2. placement of modules by exploiting circuit regularities;
3. rectilinear Steiner tree problems; and
4. point dominance problems.

In floor-planning, topological aspects of the problem are considered and geometric issues such as sizing are explored. A clustering technique for module placement, that exploits regularities in circuit structures, is being investigated. In this way natural clusters that reflect the hierarchical perspective of circuit connections can be built automatically. Problems in approximate designs for global and single layer routing using Steiner trees are being investigated. Point dominance, from computational geometry, is being applied to circuit layout problems.

University of Illinois; Prithviraj Banerjee; *Parallel Algorithms for Synthesis and Test*; (MIP-9320854 A001); \$42,000; 12 months.

This research is on developing circuit design and test algorithms that run on parallel computers. These include efficient, asynchronous, portable parallel algorithms for:

1. synthesis of combinational circuits; and
2. for test generation and fault simulation of combinational and sequential circuits.

Algorithms are being written using an environment that makes it possible to port CAD applications across a wide range of MIMD machines. In addition they are designed to allow a maximum overlap of computation and communication. The algorithms are being tested on numerous parallel platforms.

University of Illinois; Rajesh K Gupta; *CAREER: Architecture and Synthesis Techniques for Embedded Systems*; (MIP-9501615); 36 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$100,000).

This research is on synthesis techniques for micro-electronics- based "embedded systems". An embedded system is one designed for a specific functionality under stringent constraints on its timing performance and cost. Design automation techniques are being developed for use in a framework where the designer can assess tradeoffs between various embedded system architectures and designs. The main capabilities of this framework are:

1. timing analysis for both execution delay and rate constraints;
2. embeddable software and runtime system generation under timing constraints;
3. software size-performance tradeoffs;
4. cost- performance analysis of architectural alternatives, such as pre- fetch and forward hardware.

This work is exploring system partitioning and transformation techniques to build system implementations that "guarantee" constraint satisfaction while optimizing system development cost.

University of Illinois; C. L. Liu; *Computer-Aided Design of VLSI Circuits*; (MIP-9222408 A002, A003); \$109,393; 12 months.

The research is on algorithms for high level synthesis and layout of VLSI CAD designs. Algorithms for complex, large industrial type design problems are being developed. Research topics include:

1. high level synthesis with testability as an important goodness measure,
2. timing driven placement algorithms for EPGAs, and
3. channel and switchbox routing in which the effect of cross talk is to be taken into consideration.

For the first problem, the effect of register allocation and functional unit binding on the testability of the circuit is being examined. Then the scheduling step is being examined. Research on the second problem is based on the notion of a "neighborhood graph" which is used in guiding an iterative improvement algorithm that produces placements which satisfy given timing constraints. An integer programming approach is being used for the third problem, because this avoids parallel long wires that are close together in the routing solution.

Indiana University; Steven D Johnson; *Decomposing Digital-System Specifications into Interacting Sequential Processes*; (MIP-9208745 A002); \$49,988; 12 months.

This research is on developing formal techniques to decompose higher-level system specifications into interacting sequential processes. A functional algebra is used for defining formal representations and building a set of transformations for manipulating them. The notion of "interaction schemes" is the central subject. The project is composed of four activities:

1. theoretical studies of the decomposition of digital systems; especially formal derivation of control-synchronization and data-communication protocols;
2. automation of formal transformations which correctly do sequential decompositions;
3. integration of the mechanized formal system with available CAD tools, and other reasoning systems; and
4. application of the design system to meaningful examples of digital-system designs.

Purdue University; Kaushik Roy; *CAREER: Integrated Framework for Test Synthesis, Power Optimization, and Reliable Design of VLSI Circuits*; (MIP-9501869, A001); \$90,000; 36 months; (Joint support with the Microelectronic Systems

Architecture Program - Total Grant \$100,000).

This research is an integrated approach to design of VLSI circuits, with an emphasis on developing tests for thermal and electrical stress, and power management. The first goal is to use stress testing as a low-cost alternative to burn-in. Methods to synthesize tests for electrical and thermal stress based on accurate measure of signal activity are being explored. In low power design, the following are being pursued:

1. develop Monte Carlo based approaches estimation of signal and glitching activity in sequential and DSP circuits;
2. find architectural and system-level power minimization techniques;
3. investigate current switching techniques for power management; and
4. design layout algorithms for circuits with low power consumption. Sample sequential and digital signal processing circuits are being synthesized.

University of Notre Dame; Edwin H Sha; CAREER: High-Level Design Methodologies for Time-Optimal and Memory-Optimal Systems; (MIP-9501006); \$114,000; 36 months.

This research is on optimization algorithms for synthesis of multi-level loops, which occur in time and memory critical parts of scientific computing applications. The nested loops are modeled as multi-dimensional data-flow Graphs (MDFG); and algorithms taking advantage of the multi-dimensionality are being designed. By considering the multi-dimensional iteration space and the iteration body simultaneously, the transformation and optimization techniques are able to optimize throughput and memory requirement at the behavior level. Research topics include: graph transformation and optimization; data scheduling; and co-design. Polynomial-time algorithms for various graph models are being developed. This avoids exponential integer linear programming approaches.

Iowa State University; Liang-Fang Chao; RIA: Optimizing Synthesis for Periodic Real-Time Applications; (MIP-9410080 A001).

This research is on finding pipeline schedulers for the behavioral description of an iterative or recursive algorithm or loops. This extends the scheduling algorithm to more realistic resource models and to graphs with conditionals. The goal is to expose the parallelism in an iterative or recursive algorithm to provide information for optimization. Focus is on the innermost loops or iterations, which are the most time-critical for applications. The approach is to study the structure of cyclic data-flow graphs with edge

delays in order to optimize behavioral transformation, scheduling and partitioning.

Massachusetts Institute of Technology; Anantha P Chandrakasan; CAREER: Methodologies, Tools, and System Design of Low-Power Wireless Multimedia Systems; (MIP-9501995); \$124,000; 36 months.

This research is on design techniques for circuits with a power supply voltage of 0.5V. Innovations in technology, circuits, and architecture are being explored. Issues being pursued in the design of low-threshold devices are: when to shut down and efficient shutdown. Leakage currents are a major power drain in using low-threshold devices. Thus, techniques to efficiently detect and power down idling hardware are being devised. The potential of using adaptive supply voltages to exploit time-varying computational requirements is being investigated. Studies of the fundamental limits of reducing energy per computation are being developed into a general theory of DSP algorithms which can use dynamic power supply voltages. Such algorithms change computational complexity dynamically, thus saving power. Circuits which operate at supply voltages less than 0.5V are being fabricated.

Massachusetts Institute of Technology; Srinivas Devadas; NYI: Formal Methods for Hardware and Software Verification; (MIP-9258376 A004); \$62,500; 12 months.

Research is on logic and behavioral verification of VLSI circuit designs, and application of hardware verification techniques to software verification. Topics include:

1. Use of Free Binary Decision Diagrams (FBDDs) to find useful Boolean representations of circuits and efficient manipulation methods for them. Algorithms for combinational and sequential, synthesis, test and verification applications are being developed.
2. Automatic methods to verify pipelined implementations against unpipelined specifications are being explored. The methods ensure that each data transfer that takes place upon the execution of any instruction in the unpipelined circuit also occurs in the pipelined circuit. A symbolic simulation method is being developed that will efficiently verify pipelined micro-processors against instruction set specifications.
3. FBDD representations are being used to find symbolic traversal methods which allow for automatic software verification. These are also being used to debug software programs by verifying that the program satisfies correctness properties.

University of Massachusetts - Amherst; Premachandran Menon; Optimization of Multilevel Logic; (MIP-9311185 A001); \$42,055; 12 months.

This research is on the development and implementation of optimization techniques for multilevel combinational logic circuits specified at the gate level. Logical implementation techniques which have been proven effective in test pattern generation are being used for this purpose. These techniques are being applied to three

problems in logic optimization: area reduction, delay reduction, and testability enhancement. Relationships between signal values are being used to identify common subfunctions without generating expressions for them and putting them in any standard form. This approach is efficient and could substitute for Boolean factorization in logic optimization. Techniques for delay reduction involve replacing segments of long paths with shorter ones so that false paths are not created and undetectable faults are identified during delay reduction.

The effectiveness of these techniques as supplements to an existing synthesis system are being investigated.

Princeton University; Sharad Malik; *NYI: Design Automation for Embedded Systems*; (MIP-9457396 A001); \$65,000; 12 months.

High fabrication costs coupled with decreasing time-to-market for products require an increase of the programmable component of integrated circuits. Core processors embedded in a gate array provide programmability. This changes circuit design from just gates in an array, to a combination of instructions running on a core processor coupled with a gate array. Gates, as basic units of computation on silicon, are well understood. This research is focused on understanding embedded instructions as basic units of computation on silicon. Initial work is on finding methodologies for synthesizing embedded software. Problems being examined are: worst case timing analysis; power modeling and analysis; and high- quality retargetable code synthesis. Approaches include: developing timing models capable of handling features such as caches and pipelines, and determining worst-case paths; analyzing system power consumption; and formal architectural specification of specialized architectures, as well as code-generation algorithms that can take advantage of special architectures.

Princeton University; Wayne H Wolf; *Architectural Co-Synthesis for High-Performance Distributed Embedded Systems*; (MIP-9424410); \$125,475; 36 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$192,383).

Embedded computing systems must be designed to meet hard, soft, performance, cost and other constraints. This research studies the synthesis of embedded systems built from multiple processors, which may be commercial or application-specific ICs connected in a network. This methodology, co-synthesis, aims to simultaneously design the hardware and software architectures of a system. Algorithms for meeting all constraints through iteration of both the hardware architecture of the distributed computing engine and the process architecture of the application software are being investigated. The model for the research is the task graph and distributed system

graph. These are being extended based on simultaneous design the hardware and software, stronger assumptions about processes in the task graph, and simultaneous scheduling, allocation, and partitioning of processes during co-synthesis. Mathematical optimization methods, such as graph partitioning and network flow, are being used to generate effective design algorithms.

Columbia University; Steven M Nowick; *CAREER: Testability and Sequential Optimization of Asynchronous State Machines*; (MIP-9501880); \$120,000; 36 months.

This research is part of a long term effort to migrate techniques for synchronous design to asynchronous design. The two focus areas are sequential optimization and synthesis for testability. In the former, problem formulations of input, output and input/output encoding are being investigated. In each case, the problems of optimality, race-free encoding, and hazard free logic are being solved simultaneously. In the latter, synthesis of fully testable asynchronous state machines using partial or no scan is being pursued. Design tools, incorporating the algorithms for optimization and synthesis for test, are being developed.

Rensselaer Polytechnic Institute; Robert A Walker; *RIA: Combinational Optimization Algorithms for Scheduling and Module Selection*; (MIP-9211323 A001); \$5,000.

The objective of this research is to build a high-level synthesis system which maps behavioral descriptions into register-transfer level designs with guaranteed performance. The approach is to examine subtasks of the high-level synthesis problem using theoretical tools necessary to understand the synthesis problem. Synthesis subtasks in the areas of scheduling and datapath allocation are being examined. Scheduling problems being investigated include: finding good heuristics to bound the search space so that good quality schedules for large size designs can be found in a timely way; use an integer programming formulation to solve the module selection problem; and develop a powerful, integer programming based, algorithm to optimally group functions into hardware modules. Register allocation problems include: devising a register allocation algorithm which will minimize both the number of registers and register transfers in the presence of loops and conditionals while accounting for interconnect cost; formulate the module allocation problem along the same lines as the register allocation; and investigate a good ordering of the allocation subtasks so that interconnects play a major role in the design process.

Rensselaer Polytechnic Institute; Robert A Walker; *Solving the "Extended" Scheduling Problem*; (MIP-9423953); \$174,040; 24 months.

This research is on formal analysis of the scheduling problem in high-level synthesis. Problems being addressed are:

1. finding the length of the control step (clock estimation);
2. determining the type of functional unit that will perform each operation (type mapping);
3. computing the lower bound on overall schedule length; and

4. estimating the lower bound on the number of functional units.

State University of New York - Stony Brook; Michael M Green; *NYI: Improved Circuit Simulation Using Results from Circuit Theory*; (MIP-9457387 A001); \$31,250; 12 months; (Joint support with the Circuits and Signal Processing Program - Total Grant \$62,500).

This research is applying the principal investigator's previous work in the area of nonlinear circuit theory to make major enhancements to the way designers simulate analog circuits. In particular, improvements to the continuation methods of solving dc operating points of circuits are being made guaranteeing that all of a circuit's operating points will be found during a single analysis. Moreover, continuation methods are being applied to sensitivity analysis of circuits; for example, by making observations of a continuation curve, a designer could determine whether a circuit is prone to a latch up condition. Erroneous models are thought to be a major source of convergence problems and erroneous results in circuit simulation. Another enhancement to circuit simulation includes checking the accuracy of transistor models by verifying that all models satisfy passivity and the no- gain condition.

University of Rochester; Eby G Friedman; *Automated Synthesis of High Performance Clock Distribution Networks*; (MIP-9423886); \$195,000; 36 months.

This research is on automated synthesis of high speed, highly reliable clock distribution networks. A global clock signal is required in order to control synchronous operations, and it must be distributed to every register at a precise time. A four phase top-down design system for synthesizing buffered clock distribution networks is being investigated. Models of operation which include the effects of process parameter variations on timing are being developed. Scheduling clock skew is done at the behavioral level of the system. The clock distribution algorithms are being included in an integrated synthesis system.

University of North Carolina - Charlotte; Dian Zhou; *NYI: Performance-Driven VLSI Designs*; (MIP-9457402 A001); \$31,250; 12 months; (Joint support with the Systems Prototyping and Fabrications Program - Total Grant \$62,500).

Three research issues in high-performance VLSI system design are being addressed:

1. how to relate the system performance function, characterized by electrical parameters, to the geometrical parameters of the VLSI physical design;
2. how to model performance driven VLSI physical designs based on given technology and computational capability; and
3. how to characterize the fundamental computational aspects of modeled problems and develop effective algorithms for solving them.

A distributed RLC circuit model for interconnects is being designed. It considers: non-monotone circuit response, coupling effect among the signal lines, and low energy consumption. Efficient computation methods that solve time-varying Maxwell equations using the adaptive wavelet collocation method (AWCM) are being devised. The algorithms and methods are being included in a CAD

system.

because functional behavior and data values generated by modules are not considered.

Carnegie-Mellon University; Edmund M Clarke; *Automatic Verification of Finite-State Concurrent Systems in Hardware and Software*; (CCR-9217549 A002); \$37,000; 12 months; (Joint support with the Software Engineering Program - Total Grant \$148,300).

Logical errors in sequential circuit designs and communication protocols have always been an important problem. They can delay getting a new product on the market or cause the failure of some critical device that is already in use. The most widely used method for verifying such systems is based on extensive simulation and can easily miss significant errors when the number of possible states is very large. This research deals with developing an alternative approach based on a technique called temporal logic model checking. In this approach specifications are expressed in a propositional temporal logic, and sequential circuits and communication protocols are modeled as state transition systems. An efficient search procedure is used to determine automatically if the specifications are satisfied by the transition system. The technique has been used in the past to find subtle errors in a number of non-trivial designs.

University of Pittsburgh; Steven P Levitan; *Temporal Specification Verification*; (MIP-9102721 A003); (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$25,000).

This research is on verifying timing specifications for interconnection of modules in both synchronous and asynchronous digital systems. The notion of temporal behavior is being abstracted from the notion of functional behavior by focusing primarily on the control protocols of the modules and ignoring the data values computed by the modules. In this model, the interface protocols of each module are given along with the connectivity between modules. A static graph is built that describes the temporal relationships among all the external signals of all the modules. The verification process is based on a comparison between the possible behaviors of the system, represented by the graph, and the legal behaviors as represented by a set of constraints. The key constraint is that the temporal behavior of one module cannot violate the temporal constraints of another module within the system. The algorithms support multiple system states, state transitions, and checking of conditionals and loops within the protocols. This searching is tractable

University of Pittsburgh; Steven P Levitan; *Computer Aided Design of Electro-Optical Information Processing Systems*; (MIP-9421777); \$49,570; 12 months.

This research is on defining requirements for computer aided design of electro-optical information processing systems. Three steps are being pursued. First, define levels of abstraction for electro-optical systems, analogous to the behavioral, structural, electrical and physical abstraction levels of VLSI design. Second, characterize functional, physical and parametric models necessary to analyze designs at different levels of abstraction. Third, specify requirements for a multi-level simulation system.

Southern Methodist University; Sukumaran Nair; *RIA: Spectral-Based Numerical Methods for Combinational Logic Synthesis*; (MIP-9410822 A001); \$10,000.

This research is on numerical methods for digital logic synthesis based on spectral information of the logic description. The approach is to use the BDD description of the circuits to develop efficient methods for computing the spectral coefficients. Because of the numerical nature of the algorithms, design optimizations can be included in the synthesis process, rather than later. Theoretical results are being derived and used to implement a new class of synthesis algorithms that are suitable for incorporation into existing design environments. The algorithms are being tested on industry benchmarks.

Brigham Young University; Phillip J Windley; *A Distributed, Type-Based Library of Abstract Hardware Modules*; (MIP-9412581); \$130,994; 24 months.

The goal this research is to make sharing of verified hardware designs more practical. First, a model for describing and verifying type-based abstract modules, which can be used and reused in different HDL's and proof systems, is being developed. This model provides a methodological approach to verification. It is based on work in generic interpreters and micro-processor verification. Second, a model for interchange and sharing library components is being developed. This is to enable designers to effectively develop and use module libraries in a distributed manner. The module libraries being developed are publicly available so that the ideas about reuse can be tested on a large scale.

University of Utah; Erik L Brunvand, Ganesh Gopalakrishnan; *The Design of Asynchronous Circuits and Systems with Emphasis on Correctness and Proven Optimizations*; (MIP-9215878 A001); \$42,610; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$85,220).

This research merges two efforts in asynchronous circuit compilation. These are the work of Gopalakrishnan on the language hopCP and its use in verification; and the work of Brunvand on asynchronous circuit compilation. The research is:

1. enhancing the expressive power as well as the semantic clarity of concurrent hardware description languages for asynchronous circuits and systems;
2. extending the formal basis for compiling from HDL's to circuit designs;
3. formally characterizing and improving the optimizations used in asynchronous circuit compilation; and
4. studying the performance of implemented circuits with regard to a variety of parameters.

University of Utah; Ganesh Gopalakrishnan; *A Multi-Paradigm Verification System Tailored for the Design Refinement Cycle*; (MIP-9321836 A001); (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$5,000).

The design of a VLSI system involves multiple design representations; and the design must go through several iterations aimed at meeting many performance and cost constraints. Verification that the design meets constraints is necessary. This research is developing rigorous verification methods that span multiple design representations, accommodate design revisions, and provide incisive partial verification methods (e. g. verification focussed on the "corners" of the behavioral

space) that fit within designers' time budgets. These ideas are being validated by verifying the real asynchronous designs.

University of Virginia; Gabriel Robins; *NYI: New Directions in High-Performance VLSI Layout*; (MIP-9457412 A001, A002); \$72,500; 12 months.

Realistic formulations of performance-driven layout are the focus of this research. Accurate models of circuit delay are being sought. These models include technology parameters, such as capacitance, resistance, inductance, etc. The approach is to study delay-optimal trees to define an envelope of achievable routing performance. Methods for constructing near-optimal layouts are being investigated. Additionally, the routing problem is being recast as one of constructing low-delay routing graphs where cycles are allowed. This can have the advantage of designs being tolerant to certain types of open faults due to manufacturing defects or electro-migration.

Virginia Polytechnic Institute and State University; James R Armstrong, Walling R Cyre; *Rapid Development and Testing of Behavioral Models*; (MIP-9120620 A005); (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$20,000).

This research is on methods to create behavior models that accurately represent the functionality and timing of complex devices.

The work is on developing a "Modeler's Assistant" as a base for structured development of behavioral models. Specific problems being solved are: developing a process primitive set for the Modeler's Assistant, developing and evaluating performance measures for structured behavioral model development, developing a natural language interface for the Modeler's Assistant, and building into the Modeler's Assistant the capability to automatically generate tests for any behavioral model which has been constructed by the system. Behavioral models are being expressed in the high-level language, VHDL.

University of Washington; Steven M Burns; *NYI: A Design Language for Asynchronous Circuit Synthesis*; (MIP-9257987 A003); \$62,500; 12 months.

This research is on developing a unifying design language and framework in which asynchronous circuit designs can be completely specified, and in which decisions made during synthesis of the implementation can be recorded. The language is an extension of Hoare's CSP with a means to specify structural hierarchy with a refinement hierarchy superimposed upon it. The research consists of three interconnected tasks: language definition; developing tools such as a parser, flattener, view generator, handshaking expansions, and production rule sets; and application of the tools to a large design or a real time system. An algorithm for determining the maximum time separation of event occurrences in a concurrent system is being developed.

University of Washington; Carl Sechen; *Symbolic Analysis of Large Analog Circuits*; (MIP-9406470 A001); \$46,000; 12 months.

This research explores algorithms and techniques for symbolic analysis of large linear or linearized integrated circuits in the complex frequency domain. The approach is to extract transfer functions of circuits in symbolic form. In the case of large circuits, approximate symbolic network functions, in expanded or nested format, are generated. Both perturbation-based, and tree enumeration approaches are being used. The symbolic algorithms and simulators are being integrated into an analog design automation system.

University of Washington; Andrew T Yang; *NYI: Modeling and Simulation of Advanced Microelectronics and Optoelectronic Circuits and Systems*; (MIP-9257279 A003); \$62,500; 12 months.

This research is on modeling and simulation of mixed analog/digital circuit designs. Topics include:

1. Hierarchical modeling and simulation of tightly coupled, mixed digital circuits;
2. Simulation of high-performance circuits with complex nonlinear and electromagnetic effects;
3. Modeling and simulation of optoelectronic integrated circuits, with the focus on automatic model generation;
4. Development of a technique for fast power waveform simulation based on analytic digital macromodeling, which makes the method well suited to simulating circuit netlists derived from layout extraction;
5. Fast simulation of interconnect problems using asymptotic waveform evaluation so as to obtain efficiency and accuracy in approximating nonlinear portions of the circuit.

Algorithms for solving these problems, including methods for simulating designs with over 100,000 transistors, are being developed.

Manufacturing Test

Stanford University; Edward J McCluskey; *Research on Reliable Computers*; (MIP-9107760 A002); \$150,001; 12 months.

This research is developing techniques for reducing the occurrence of run-time errors in circuits and systems. Emphasis is on preventing the introduction of faults into the system through verification and synthesis techniques, and on improving the detection and diagnosis of faults causing run-time errors so that the faults can be removed from the system. Topics being pursued are:

1. Developing synthesis methods that automatically synthesize a synchronous, register-transfer level (RTL) hardware specification from a behavioral VHDL language specification;
2. determining what coverage of multiple stuck-at faults or bridging faults can be expected or guaranteed by a test for delay faults in an arbitrary circuits;
3. investigating methods for utilizing the use of output waveform characteristics in delay testing; and
4. finding fault models for intermittent failures, and methods to detect their presence.

University of California-Santa Barbara; Kwang-Ting Cheng; *Reachability Computation Using the Extended Finite State Machine Model and its Application to Automatic Test Generation*; (MIP-9503651); \$129,957; 36 months; (Joint support with the

Microelectronic Systems Architecture Program - Total Grant \$201,844).

This research addresses computation of reachable states for VLSI designs described in high-level languages and the application of these results to automatic test pattern generation (ATPG). Methods to compute symbolically the set of states reachable from an initial state are being explored. An extended finite state machine (ESFM) model, which can represent communication protocols and hardware behavior, is being used. A key idea is that the model allows use of arbitrary state variables, such as boolean and arithmetic, which will provide efficiency in computation. Methods and prototype tools to convert automatically a design in VHDL to an ESFM are being developed. These methods and techniques are being applied to ATPG for sequential circuit test.

University of California-Santa Cruz; Frankie J Ferguson; *PYI: Hierarchal Test Pattern Generation for Manufacturing Defects*; (MIP-9158491 A005, A006); \$72,726; 12 months.

The focus is on developing cost-effective testing methodologies that detect more defective ICs than current methods. There are two approaches to manufacturing test. The use of high-level fault models reduces test generation costs, but furnish lower quality tests. The use of low-level fault models increases the quality of circuits that have passed the tests, but causes testing costs to mushroom. This research integrates these two techniques so that tests can be generated that detect virtually all plausible manufacturing defects without excessive

automatic test pattern generation costs. The approach is to develop a software tool, "Carafe" (circuit and realistic fault detector), which determines the most likely faults to occur in a CMOS circuit. It exploits the hierarchical nature of VLSI circuit designs, making fault extraction faster and more memory efficient. Also it supports multi-level metal CMOS technologies.

University of California-Santa Cruz; Tracy Larrabee; *PYI: Sequential and Combinational Test Pattern Generation for Realistic Faults Using Boolean Satisfiability*; (MIP-9158490 A006); \$62,500, 12 months.

This research is based on an automatic test pattern generation (ATPG) system (called Nemesis) that generates a test pattern for a given fault by first constructing a formula representing all possible tests for the fault, and then applying a Boolean satisfiability algorithm to the resulting formula. This method separates the formula extraction from the formula satisfaction thus providing flexibility and generality. A testing system, based on Nemesis, that will generate tests detecting all realistic manufacturing defects in both combinational and sequential ICs is being developed.

University of Southern California; Sandeep K Gupta; *CAREER: Tool for At-Speed Robust Path Delay Testing*; (MIP-9502300); \$137,903; 36 months.

This research is on path delay testing. Tools and techniques to enable the development and evaluation of practical at-speed delay tests are being investigated. Test generators, fault simulators, and techniques for the design of design-for-testability (DFT) and built-in self-test (BIST) circuitry are being developed. Special attention is being paid to robust delay testing of the most critical circuit paths. Software tools which provide a practical testing solution for large circuits are being developed.

Georgia Institute of Technology; Abhijit Chatterjee; *CAREER: Testing of Digital Circuits by Signal Waveform Analysis*; (MIP-9502575); \$109,121; 36 months.

This research is developing a new approach to manufacturing test of VLSI chips. Digital signals are treated as analog waveforms, so that logic events at internal nodes can be understood. To do this, rise-fall and glitching characteristics of output signals are being studied using time and frequency domain analog signal analysis techniques. Using this approach:

1. failure methods are being determined;
2. test generation algorithms for stuck-at faults in combinational and sequential circuits are being

- investigated;
3. methods to detect marginal chips are being explored;
4. very rapid wave pipelined testing is being developed. A signal waveform analysis system is being designed.

Southern Illinois University at Carbondale; Spyros Tragoudas; *RIA: Built-In Test Pattern Generation Methods*; (MIP-9409905 A001); \$10,000.

This research is on the design of test pattern generators for combinational circuits. Various schemes for LSFR/SR pseudorandom test pattern generators are being examined. The focus is on finding high quality ATPG algorithms that require minimal hardware overhead. Several families of not-necessarily primitive, characteristic polynomials are being investigated with the objective of guaranteeing that test patterns will be applied randomly to all outputs. In addition, various partial scan schemes that follow the structural approach are being studied.

University of Iowa; Irith Pomeranz, Sudhakar M Reddy; *High-Quality Tests for Combinational Circuits*; (MIP-9220549 A001); \$78,896; 12 months.

This research is on finding procedures to derive compact or small test sets that cover a comprehensive set of modeled faults, required to achieve high reliability of manufactured VLSI chips. A fault model that allows uniform representation of various fault models is being explored as the basis for generating small, yet comprehensive test sets. A set of tools to deal with different aspects of testing quality for combinational and fully scanned sequential circuits is being built. These are:

1. generation of small test sets for a comprehensive set of faults, including efficient treatment of path delay faults;
2. design for testability to allow faults undetectable in the original to be detected in a modified circuit; and
3. built in test pattern generation based on the test sets produced when stored pattern tests cannot be used.

University of Iowa; Irith Pomeranz; *NYI: A New Search Strategy for Design Automation*; (MIP-9357581 A002); \$52,650; 12 months.

A method for solving CAD design and test problems, which is especially suitable for solving large problems because it does not deteriorate as circuit size is increased, is being explored. The approach is to scale down the problem while retaining all of its details; find high quality solutions for a sequence of small circuits; develop rules for solving these problems; and then scale up these miniature solutions into a solution for the large problem. Since the miniature solutions are optimal, the full- sized solutions are expected to have very high quality.

Research problems being explored are:

1. test generation for various fault models on designs given at the gate and higher levels;
2. built-in-self-test methods; and
3. testing of synchronous sequential circuits that require two pattern tests.

Michigan State University; Chin-Long Wey; *Efficient Testing Paradigms and Diagnosable Design Methodologies for Analog Integrated Circuits and Systems*; (MIP-9321255 A001, A002); \$75,000; 12 months.

This research is on test paradigms for high-frequency, linear and nonlinear mixed signal circuits and systems. Issues in BIST (Built-In Self-Test) and design for testability are being addressed. Fault models for hard faults (open or short

circuits) and for soft faults (deviations from nominal component values) are being developed. For BIST design, a high speed, low power, current mode copier is employed as the sample/hold circuit to achieve 10 ns/sample at 0.1% accuracy under 3.3V power supply. Both CMOS and BiCMOS technology are being used.

Design for testability work focuses on properly selecting component values and topological structure to increase testability and diagnosability. The problem of determining component values is formulated as an optimization problem which includes testability strategy as a parameter. The new methods are being tested on realistic circuits.

University of Michigan; John P Hayes; *Functional Testing of Complex Digital Systems*; (MIP-9503463); \$129,185; 36 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$202,000).

This research is on techniques for testing VLSI circuits directly at functional levels, while ensuring that standard testing quality goals are met. A high-level fault model called the induced functional fault (IF) model is being investigated. With this model, an IF fault class can be coupled to a physical fault class so that detection of IF faults ensures detection of physical faults. Advantages of this approach are easy derivation of near- minimal test sets, and reuse of tests in other circuits. The construction, use and reuse of precomputed tests for sequential circuits are being explored. In this regard, the application of balance and functional properties to simplifying test generation and improving fault coverage are being studied. Built-in self- testing methods that facilitate the implementation of functional testing methods and the detection and elimination of design errors are also being explored.

University of Minnesota; Bapiraju Vinnakota, Ramesh Harjani; *Testing Analog and Mixed Signal Integrated Circuits*; (MIP-9501499); \$273,648; 36 months.

This research is on analog circuit observer block (ACOB), structures for testing analog and mixed-signal circuits. Design styles for analog and mixed-signal ICs and their influences on testing the products are being studied. ACOBs for circuits of practical importance, both discrete and continuous time, are being investigated. A second exploration is into fault models for analog circuits which quantitatively link defects to variations in the functional parameters, which are traditionally verified in analog circuit test. These models are being used to validate the effectiveness of defect-oriented testing and design-for-test techniques, including ACOBs. The algorithms are being verified on practical circuits. Software for analog and mixed-signal test is being developed.

University of Minnesota; Bapiraju Vinnakota; *CAREER: Low-Cost Test Techniques*; (MIP-9502240); \$115,000; 36 months.

This research is an investigation of techniques for efficient and effective testing of large sequential VLSI circuits. Two avenues are being pursued. First is the development of self-initializing memory elements and test time reduction techniques for circuits employing scan design. Second is development of new functional test techniques for large sequential circuits. The approach is to model large circuits as a network of finite state machines. From this model, algorithms for test generation are being developed. The algorithms are being optimized to provide good coverage for defects not modeled by stuck-at faults. Software instantiations of the test algorithms are being built.

Princeton University; Niraj K Jha; *High-Level Synthesis for Hierarchical Testability*; (MIP-9319269 A001); \$77,000; 12 months.

This research is concerned with finding efficient hierarchical testability techniques for controller-data path systems. The approach is to start with module level test sets, derived for any suitable fault model, and use high level synthesis to ensure that these test sets can be combined into a system level test set which can provide complete test coverage of all the embedded modules. The aim is to invent algorithms that reduce test generation and application times, yet obtain complete, or nearly complete, system level test

coverage with little or no area and delay overhead. The testability techniques are being embedded into algorithms for scheduling and allocation. Also being explored are methods for performing synthesis with both low power and testability as design criteria.

Texas A & M University; Dhiraj K Pradhan, Duncan M H Walker, Wolfgang Kunz; *Novel Methods in Computer-Aided Circuit Design and Testing Using Recursive Learning*; (MIP-9406946 A001); \$108,969; 12 months.

This research investigates a methodology, "Recursive Learning" (RL), which provides a conceptual solution to the Boolean satisfiability problem by determining complex implications within a logic circuit or expression. Applications within a logic circuit or expression. Applications of recursive learning are being made to:

1. Logic verification,
 2. Multi-level verification,
 3. Random pattern testable synthesis, and
 4. multi-level optimization using Galois logic.
- Algorithms and tools based on these results are being developed.

University of Wisconsin; Charles R Kime; *Built-In Self-Test for Random Pattern Resistant Faults*; (MIP-9319742 A001); \$65,000; 12 months.

This research is on techniques for testing circuits having random pattern resistant faults. The focus is on methods for synthesizing random pattern testable circuits. The synthesis algorithms being developed include algorithms for design of random pattern testable two level and multi-level combinational logic; extensions to some classes of sequential logic are also being investigated. Two methods being pursued are weighted cellular automaton, and fixed bit biased pseudo-random pattern generator.

Simulation

University of Texas at Austin; Larry T Pileggi; *PYI: CAD Tools for New Circuit Technologies*; (MIP-9157363 A005); \$62,500; 12 months.

The research is on developing simulation capabilities at the system level, which are as powerful as those at the chip level. To this end, tools which are applicable at the integrated circuit, packaging, module, and board levels without loss of generality are being developed. Specific areas being investigated are: fast extraction techniques which trade off some accuracy for efficiency; and simulation models and techniques to enable topdown design. Timing analysis of boards or multi-chip modules requires development of tools for the extraction, characterization and simulation of transmission lines at the system simulation level. Toward this end interconnect macro- models, similar to those for on-chip RC interconnect circuits, which are

compatible with system-level design automation tools, are being explored.

Other

California Institute of Technology; Michelle Effros; *Code Clustering for Universal Image Coding and Other Implications*; (MIP-9501977); \$38,119; 36 months; (Joint support with the Systems Prototyping and Fabrications Program, Circuits and Signal Processing Program, Microelectronic Systems Architecture Program - Total Grant \$134,628).

This project seeks to develop reasonable complexity, source-independent coding algorithms, which are crucial to the design of robust systems for image coding and mobile communications. In these applications the statistics of the source and channel in operation are typically unknown a priori, and the performance of the coding strategy employed is sensitive to those unknown characteristics.

The two-stage approach developed in the source coding literature is employed. The literature demonstrates that in general one should quantize the space of possible codes. Some of the rate should be spent on describing which code, in a family of codes, should be used on the source in operation. Specific projects include the development of a universal DCT code compatible with JPEG and MPEG image and video standards, a universal KLT code, a universal wavelet packet code, and a universal channel code.

The main objectives of the education plan are to develop and maintain an exciting atmosphere for active learning for undergraduate and graduate students through innovative programs that encourage the maximum possible exchange between students, faculty, and individuals from local industry.

University of California-Berkeley; Ljiljana Trajkovic; *VPW: Homotopy Methods for Analysis and Simulation of Electronic Circuits*; (DGE-9550153); \$15,000; 15 months; (Joint support with the Visiting Professorships for Women, and the Western Europe Program - Total Grant \$171,397).

Finding dc operating points, steady state, and transient responses of electronic circuits are essential tasks in electrical circuit simulation and involve nonlinear differential/algebraic equations. Traditional methods for solving such systems of equations often fail, are difficult to converge, and, often cannot find all the solutions. Dr. Trajkovic will investigate the application of homotopy methods to solving nonlinear equations describing power electronic and switching

circuits and power systems flow equations, that traditionally pose simulation difficulties. Experiments with homotopies may lead to the development of better circuit simulation tools and to better understanding of relationships between homotopy methods and the behavior of nonlinear circuits. Dr. Trajkovic will spend three months with the Electrical Engineering Department at the University of Wuppertal, Germany where she will collaborate with Professor Wolfgang Mathis and his students. Their "Theory and Computer Aided Design of Electronic Circuits" group actively pursues research in circuit theory and simulation. This will be a continuation of her research activities at Berkeley, with emphasis on theoretical work based on experiments with various homotopies.

Interactive activities include: teaching a graduate course in "Advanced Circuit Theory" that will introduce students to analytical results that provide insights and understanding of complex behavior of electronic circuits; holding a workshop/seminar series on "Application of Homotopy Methods in Solving Engineering Problems," with special emphasis on circuit simulation and solving equations describing nonlinear circuits; and organizing informal workshops where female students make presentations and get advice about their research and career goals.

University of Colorado; Michael Lightner; *Workshop on the CAD Needs for Supporting System Design for the Next Ten Years*; Boulder, Colorado, March 1995; (MIP-9504395, A001); \$18,598; 6 months.

This workshop is to develop a seminal statement about research topics needed in developing a science of tools for electronic systems design. It brings together academic and industrial researchers with experience in VLSI chip, electronic system, and other design automation fields to integrate their varied points of view. The workshop facilitates the integration of a wide range of design automation research ideas, and their integration into a new science for systems design.

University of Delaware; Phillip Christie; *Statistical Analysis of Interconnect-Limited Systems*; (MIP-9414187, A001); \$32,499; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$57,499).

The growth in size and complexity of modern computers is matched by the need for more accurate models of their internal wiring structure. The task of modeling this complexity is ideally suited to a recent development in statistical physics called the renormalization group. It has the power to attack the wiring optimization problem because it takes this one very difficult problem and breaks it up into a very large number of much simpler problems. Rather than attempting to attack the entire range of partitioning and placement requirements from the transistor level to the back-plane, an optimization procedure need only be found for a highly restricted wire length range. Once this has been derived, the procedure is renormalized for all other length scales in a manner which generates a system-wide solution. This project will investigate how such an approach may be used to develop much better estimates of interconnect-limited system performance and how this theory leads to new optimization algorithms. In contrast to other theories the predictions made by this technique will grow ever more accurate as the systems grow larger and more complex.

University of Kansas; Frank M Brown; *A Theory of Design Change*; (MIP-9526532); \$20,000; 12 months; (Joint support with the Numeric, Symbolic & Geometric Computation Program - Total Grant \$34,709).

The mental process of predicting the reasonable consequences of a change in the design of an artifact is an important step in the design process. This research addresses issues in developing a theory of design change. The premise is that the incremental changing of the geometric design of an artifact may be formally explained by logical and mathematical theories, similar to those in automated planning and simulation systems which are used in many fields. Topics being investigated are representation of designs, designer intentions, and constraints. A computational model of the deductive process of changing a design is being developed. The model uses frame laws and default constraints expressed in the modal quantificational logic Z. Other work includes exploring problems of automated deduction of facts in three dimensional geometry, consistency analysis, and developing a taxonomy of appropriate engineering design actions and intents.

State University of New York - Stony Brook; Armen H

Zemanian; *VLSI Interconnection Networks*; (MIP-9423732); \$100,000; 24 months.

This research is on techniques for analysis of VLSI interconnect nets. Methods for predicting capacitance effects in VLSI circuitry, especially at rounded corners, are being investigated using asymptotic expansions. Secondly, methods to determine the voltage-current regime of nonlinear networks on a chip are being assessed. The approach is to model each interconnection network as a cascade of three terminal networks connected together at their ends. Each cascade can be described mathematically by a nonlinear operator. Computationally efficient analysis techniques are being developed.

Carnegie-Mellon University; Randal E Bryant; *1995 Dagstuhl Workshop on Binary Decision Diagrams, Germany*; (MIP-9503339); \$7,500; 6 months; (Joint support with the Programming Languages and Compilers Program - Total Grant \$15,000).

Binary Decision Diagrams (BDD's) have found widespread use in synthesis, formal verification and testing of digital circuits. This success in the EDA research area has spawned research efforts on a number of fronts, including theoretical studies of algorithms and complexity, applications to such areas as artificial intelligence and logic programming, and extensions beyond Boolean functions to represent matrices, Markov systems, and multi-variate polynomials.

This workshop (February 13-17, 1995) is the first ever held specifically on this important technology.

Carnegie-Mellon University; Roy A Maxion, Andrzej J Strojwas, David L Banks; *Discovering Information in Large, High-Dimensional Databases*; (IRI-9224544 A002); \$10,000; 12 months; (Joint support with the Knowledge Models and Cognitive Systems, the Database & Expert Systems, the Experimental Systems Program, and the Statistics Program - Total Grant \$160,000).

Overcoming the "curse of dimensionality" is a vital problem for any complex manufacturing industry, such as VLSI production, in which hundreds of variables must be precisely controlled in order to achieve high-quality yield. This research addresses both theoretical and practical concerns. On the theoretical end, statisticians have recently proposed a number of compelling new ideas for high-dimensional, nonparametric regression (e. g. ACE, AVAS, LOESS, PPR, MARS, RPR and several other algorithms). These ideas are largely untested, and little is known about their comparative performance in realistic situations. To remedy this, a large-scale simulation experiment is performed that employs statistical design to evaluate the effects of sample size, dimensionality, signal-to-noise ratio, and various kinds of underlying functions on the integrated mean squared error of the fitted model. The results of the study are examined in an analysis of variance, leading to clear conclusions as to the circumstances under which each of the proposed methods is most valuable. On the practical side, this research applies the methodology studied in the simulation experiment to VLSI production data as micro-modeled by the PREDITOR software, which is widely

used in industry to calculate from physical principles the actual result of each step in the production of a VLSI circuit wafer.

Brown University; Harvey F Silverman; *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A002, A003, A004); \$12,060; 24 months; (Joint support with the Systems Prototyping and Fabrications Program, the Experimental Systems Program, the Circuits and Signal Processing Program, and the Microelectronic Systems Architecture Program - Total Grant \$674,601).

This is a joint project between Brown and Rutgers to build a large microphone array with associated processors for beam forming.

Applications are direction finding, echo cancellation, and speaker differentiation in teleconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

Microelectronic Systems Architecture

Dr. Lionel M. Ni, Program Director
(703) 306-1936 lni@note.nsf.gov

Program Description

The Microelectronic Systems Architecture (MSA) program supports basic research on innovative design of computer systems at the physical and the system level to achieve high system performance. It encourages studies on the impact of new hardware and software technologies, as well as the impact of new applications and algorithms on computer system architectures. The style of architectural research employed includes theoretical and analytical studies, simulations and limited proof-of-concept prototyping. The program emphasis is on physically realizable systems.

TECHNOLOGY-DRIVEN ARCHITECTURES

The program supports architectural research which explores the capabilities and limitations of current and future hardware and software technologies. The objective in these studies is to better understand and to extend the performance, programmability, applications span, and reliability of microelectronic systems. Typical issues which are addressed include:

- * Methodologies for system design that map high-level abstractions and system specifications to low-level physical implementations while considering the design tradeoffs of chip area, power consumption, clock rate, packaging, cost and programmability.
- * Design of general-purpose and special-purpose computers, such as superscalar processors, parallel processors, distributed and real-time systems. The design issues may include cache and high performance memory systems, multi-threading, interconnection strategies, pipelines, networking and I/O systems, co-processor architectures, etc.
- * Studies of system programmability, architectural support for programming languages and system software, compiler techniques to exploit and to enhance system architectural features.
- * Studies of both hardware and software strategies to enhance the reliability, availability, and fault tolerance of microelectronic systems.

APPLICATION-DRIVEN ARCHITECTURES

The program supports the design of special-purpose computers for applications that can better utilize emerging microelectronic technologies in a cost-effective manner. Projects focusing on the design and development of application-driven computing systems must involve innovative architectural research and a cost/performance analysis and evaluation of the resulting design. Primary emphasis is placed on obtaining new architectural and design knowledge and evaluating their effectiveness. A secondary emphasis is placed on studies that can provide a better understanding of the problem solving methods using microelectronic technologies. Typical issues which are addressed include:

- * Requirement specification, analysis, decomposition of problems and mapping of problem subcomponents onto functional building blocks, and analysis of the cost-performance trade-offs;
- * The design and evaluation of special-purpose computers and their required software for applications whose requirements, such as performance, memory size and physical size, cannot be met by available general purpose computers (for example, speech processing, graphics, simulation, image processing, signal processing, artificial intelligence and neural networks).

Initiatives and Opportunities

The Microelectronic Systems Architecture program is actively soliciting proposals particularly in the areas which relate to the Strategic Implementation plan for the "American Age of Information," developed by the Committee on Information and Communications of the National Science and Technology Council.

Possible research includes:

- * Design and evaluation of high performance memory systems, including I/O, for superscalar and parallel machines.

See the Dear Colleague Letter NSF 94-75 and the NSF Workshop on High Performance Memory Systems Final Report (Report No. TR-93-35, Computer Science Dept., University of Virginia, June 1993).

- * Innovative application-specific machine architectures that can tackle grand challenge problems, e.g. biotechnology and environmental studies, etc.
- * Design and evaluation of innovative microelectronic systems using new device and packaging technologies such as optoelectronics, optical interconnects, VLSI, GaAs, MCM packaging, and analog-digital devices.
- * Performance evaluation of microelectronic systems using a combination of analytical modeling, simulation, benchmarking and measurements on such systems.
- * Experimental research that deals with building of small-scale, proof-of-concept prototypes, simulation or emulation of new system designs using software or FPGA emulators.

Awards

Technology Driven Architecture

University of Arizona; Ahmed Louri; *Design of 3-D Optical Interconnects for High-Density Chip-to-Chip and Board-to-Board Interconnections*; (MIP-9310082 A002); \$103,421; 12 months.

As device speeds rise, communication rather than device speed becomes the limiting factor in performance and cost of high-performance computing systems. Optical interconnects offer the potential for high-speed, scalable communications that can keep up with progress in devices. This research explores the application of free-space optics to high-density interconnects that are capable of bandwidth, interconnect density, and error rates far beyond the capabilities of current electrical interconnects and backplanes. The approach consists of:

1. developing suitable optical network topologies;
2. identifying optical implementation techniques and required devices;
3. identifying interface components;
4. developing modeling and simulation techniques for evaluating the performance of the chosen topologies and implementations; and
5. physically implementing and measuring some of the resulting networks.

University of California-Davis; Matthew Farrens; *NYI: High Performance Single Chip VLSI Processors*; (MIP-9257259 A004); \$25,000; 12 months.

The research is to investigate various configurations of decoupled architectures and to extend the concept into the field of parallel processing. It is anticipated that, with several decoupled processors communicating via architectural queues called Multiple Instruction Stream Computer (MISC), it will function as a type of dynamic superscalar processor, providing significant performance gains. The MISC architecture uses multiple asynchronous processing elements to separate a program into streams that can be executed in parallel, and integrates a conflict-free message passing system into the lowest level of the processor design to facilitate low latency intra-MISC communication. This approach allows for increased machine parallelism with minimal code expansion, and provides an alternative approach to single instruction stream multi-issue machines such as superscalars and VLIWs. The relationship between optimal processor configuration and transistor count is also being investigated. The goal is to define the design points at which a change to multiple processors becomes advantageous.

University of California-Los Angeles; Milos D Ercegovac; *Arithmetic Algorithms and Structures for Low-Power Systems*; (MIP-9314172 A001); \$96,770; 12 months.

The objective of the proposed research is the development of numerical computing systems which operate with a small amount of electrical power. It concentrates on arithmetic structures and algorithms and models power consumption by the number of signal transitions. These low-power structures are essential for the development of future computing systems. The proposed research will further the understanding of the effect of number representation, algorithm, and implementation on the power dissipation of numerical computation structures. Moreover, methods and technique for the design of low-power structures will be developed, as well as low-power designs for specific operations, such as adders, comparators, multipliers, and dividers.

University of California-San Diego; Chung-Kuan Cheng; *Research on Circuit Partitioning*; (MIP-9315794 A002); \$24,923; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$49,846).

This project is an investigation into a new generation of hierarchical partitioning methods motivated by the need to:

1. cope with high circuit complexity,
2. improve a system's performance under I/O pin count constraints, and
3. control intermodule delay in order to optimize system performance.

This research focuses on extending previous partitioning research by adopting different circuit models including petri nets, data flows, and state machines, by improving the efficiency and effectiveness of the methods, and by deriving theoretical results on variations of partitioning formulations. The PI plans to apply partitioning methods to netlist mappings for various hardware emulation machines and to find potential applications of these methods to VLSI design problems.

University of California-Santa Barbara; Kwang-Ting Cheng; *Reachability Computation Using the Extended Finite State Machine Model and its Application to Automatic Test Generation*; (MIP-9503651); \$71,887; 36 months; (Joint support with the Design, Tools and Test Program - Total Grant \$201,844).

This research addresses computation of reachable states for VLSI designs described in high-level languages and the application of these results to automatic test pattern generation (ATPG). Methods to compute symbolically the set of states reachable from an initial state are being explored. An extended finite state machine (ESFM) model, which can represent communication protocols and hardware behavior, is being used. A key idea is that the model allows use of arbitrary

state variables, such as boolean and arithmetic, which will provide efficiency in computation. Methods and prototype tools to convert automatically a design in VHDL to an ESFM are being developed. These methods and techniques are being applied to ATPG for sequential circuit test.

University of California-Santa Cruz; Anujan Varma; *NYI: High-Speed Interconnection and Switching Technologies*; (MIP-9257103 A004, A005); \$67,500; 12 months.

This research aims to develop architectures and protocols for high-speed interconnection within computer systems. Typical applications to be considered include interconnection of processor subsystems among themselves for multiprocessor configurations, and the interconnection of processors to I/O subsystems. The focus is on high-speed crossbar switches as the interconnection means. The structure of these switches as well as mechanisms for connection setup, routing, and flow-control across multiple cascaded switches are being studied. Photonic switch architectures obtained by combining dimensions of switching, such as wavelength and space, are under investigation.

University of Southern California; Massoud Pedram; *NYI: Low Power VLSI Design*; (MIP-9457392 A001); \$50,000; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$100,000).

This research investigates modeling and estimation of power consumption as well as techniques for minimizing power at the various levels of design abstraction (layout, logic, register- transfer and behavioral levels). Principles and methods to guide the design of power efficient electronic systems are being explored; and the impact of availability of low-power design techniques on chip, module, and system level designs is being assessed. Topics being investigated include: spatio-temporal power estimation; state assignment for low power; power dissipation in boolean networks; common subexpression extraction; and FPGA synthesis for low power.

University of Delaware; Phillip Christie; *Statistical Analysis of Interconnect-Limited Systems*; (MIP-9414187, A001); \$25,000; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$57,499).

The growth in size and complexity of modern computers is matched by the need for more accurate models of their internal wiring structure. The task of modeling this complexity is ideally suited to a recent development in statistical physics called the renormalization group. It has the power to attack the

wiring optimization problem because it takes this one very difficult problem and breaks it up into a very large number of much simpler problems. Rather than attempting to attack the entire range of partitioning and placement requirements from the transistor level to the back-plane, an optimization procedure need only be found for a highly restricted wire length range. Once this has been derived, the procedure is renormalized for all other length scales in a manner which generates a system-wide solution. This project will investigate how such an approach may be used to develop much better estimates of interconnect-limited system performance and how this theory leads to new optimization algorithms. In contrast to other theories the predictions made by this technique will grow ever more accurate as the systems grow larger and more complex.

Georgia Institute of Technology; Umakishore Ramachandran; *PYI: Architectural Issues in Parallel and Distributed Computing*; (MIP-9058430 A005); \$37,500.

Beehive is a project that investigates the software and hardware issues in the design of scalable shared memory multiprocessors. The architecture is designed to support a form of weakly consistent memory model in a cache-based multiprocessor environment. The novel aspects of the architecture are reader-initiated cache coherence, latency tolerance through buffered writes, elimination of false-sharing, and efficient support for synchronization via the caches. The architectural features are designed to be compatible with an interconnection network. Research issues include: Compile-time and runtime issues in the exploration of the weak memory model and the architectural features, novel simulation techniques for the evaluation of such complex parallel systems, and exploration of optical interconnects for such scalable architectures.

Georgia Institute of Technology; Ellen W Zegura; *CAREER: A Systematic Approach to the Design of Cost-Effective, High Performance Switching Architectures*; (MIP-9502669); \$131,479; 36 months.

This research aims to answer a set of critical questions dealing with communication systems by way of the development and use of a simulation and implementation modeling environment for virtual circuit switches. The first question concerns systematic, standardized performance comparisons. Using relevant traffic models and trace data from actual applications, this study is standardizing performance measurements, eliminating the apples-to-oranges comparisons that result when considering performance results from different research groups. The second question concerns cost and implementation modeling. This work is integrating the design of system architecture and implementation. The particular focus is on electronic implementation and multichip module packaging technology. The third question concerns the role of architecture in providing quality of service guarantees. The usual abstraction of the switch as a queue hides too much detail about internal architecture that is relevant to

traffic behavior when traversing a switch. This work is characterizing the effect on traffic of different switch architectures, and developing switch mechanisms to work in concert with other mechanisms in providing quality of service guarantees.

The educational part of the project is using the simulation and implementation modeling environment as a key component of a graduate course in High Speed Switching Networks. In addition, portions of the research appear in two undergraduate courses, one on VLSI Design and the other on Modeling. Other education efforts include significant mentoring activity and ongoing improvement in teaching.

University of Illinois; Andrew A Chien; *High-Performance, Adaptive Routing in Multiprocessor Networks*; (MIP-9223732 A002); \$100,000; 12 months.

The selection of good routing algorithms in multicomputer networks is necessary to prevent deadlock, avoid hot spots, and maximize performance. This project establishes a framework for selecting routing algorithms, which consists of four parts:

1. selection of deadlock-free algorithms with varying degrees of adaptiveness;
2. a set of analytic models of router speeds based on switch and buffer delay measurements;
3. an empirical study of network traffic patterns to determine the need for adaptiveness in routing; and
4. a set of techniques for extending routers to meet requirements such as fault tolerance and in-order message delivery.

University of Illinois; Rajesh K Gupta; *Architecture and Synthesis Techniques for Embedded Systems*; (MIP-9501615); \$95,000; 36 months; (Joint support with the Design, Tools, and Test Program - Total Grant \$100,000).

This research is on synthesis techniques for micro-electronics-based "embedded systems". An embedded system is one designed for a specific functionality under stringent constraints on its timing performance and cost. Design automation techniques are being developed for use in a framework where the designer can assess tradeoffs between various embedded system architectures and designs. The main capabilities of this framework are:

1. timing analysis for both execution delay and rate constraints;
2. embeddable software and runtime system generation under timing constraints;
3. software size-performance tradeoffs;
4. cost- performance analysis of architectural alternatives, such as pre-fetch and forward hardware.

This work is exploring system partitioning and transformation techniques to build system implementations that "guarantee" constraint satisfaction while optimizing system development cost.

University of Illinois; Benjamin W Wah; *Architecture Specific Resource Management Via Intelligent Compilation and Strategy Learning*; (MIP-9218715 A004, A005); \$103,095; 12 months.

This research targets efficient distributed computing through intelligent scheduling of application programs. The approach has three components: compiler development, measurement of system loads, and automated learning of optimal scheduling policies. Compilers are modified to extract control and data dependencies from applications programs, emit performance monitoring code, and allow partitioning into processes with predictable resource requirements. A neural network model is being developed to characterize system loads based on ready list lengths, I/O traffic, and network congestion. Finally, an automated learning system will tune scheduling policies to balance system loads using the predicted and measured application program requirements.

Purdue University; Jose A Fortes; *Data Distribution Independent Parallel Processing*; (MIP-9500673, A001); \$112,919; 12 months.

This project is exploring the possibility of a parallel programming paradigm that is data-distribution independent (DDI) in the sense that the user would not be required to program or even invoke data communication routines (hereon called modules). The need for data redistribution would either be eliminated or transparent to the user. The emphasis this work is on the systematic design of computational modules so that either there is no need to redistribute input data or, when this cannot be achieved, the cost of (automatic) redistribution is minimized. In this context, source-to-source program transformations, called modular mappings, and properties that allow commutative parallel processing are being explored as techniques and concepts that enable DDI computation. The extent to which a DDI paradigm could replace existing approaches, complement them or merely apply to special application domains is unclear and this is one of the issues under investigation. In addition, hardware and architecture features that support DDI computation on both general and special purpose parallel processors are being investigated. The experimental validation is being done in the context of three application areas and program implementations on a fine-grain distributed memory SIMD machine (Maspar MP-1 with 16 thousand processors) and a coarse-grain distributed memory MIMD machine (Intel Paragon with 140 processors). The areas of interest are dense linear algebra, sparse linear algebra and symbolic compute algebra which can be applied to numerous scientific and engineering computing problems. For these three areas, DDI modules are being developed as well as entire programs built of several modules. Performance comparisons are being made between DDI implementation and their non-DDI counterparts.

Purdue University; Kaushik Roy; *CAREER: Integrated*

Framework for Test Synthesis, Power Optimization, and Reliable Design of VLSI Circuits; (MIP-9501869, A001); \$10,000; 36 months; (Joint support with the Design, Tools and Test Program - Total Grant \$100,000).

This research is an integrated approach to design of VLSI circuits, with an emphasis on developing tests for thermal and electrical stress, and power management. The first goal is to use stress testing as a low-cost alternative to burn-in. Methods to synthesize tests for electrical and thermal stress based on accurate measure of signal activity are being explored. In low power design, the following are being pursued:

1. develop Monte Carlo based approaches estimation of signal and glitching activity in sequential and DSP circuits;
2. find architectural and system-level power minimization techniques;
3. investigate current switching techniques for power management; and
4. design layout algorithms for circuits with low power consumption.

Sample sequential and digital signal processing circuits are being synthesized.

University of Notre Dame; Peter M Kogge; Architectural Techniques for Inherently Lower Power Computers; (MIP-9503682); \$155,000; 36 months.

This project focuses on the problem of power dissipation, primarily on the fundamental issues of what are the real sources of power dissipation in CMOS microprocessors, and what can be done to minimize it at a more global level. The goal is development and demonstration of techniques that would support Power-efficient Instruction Set Architectures (PISA). For this, there are four major tasks:

1. Developing technology independent models and metrics for power dissipation in CMOS logic.
2. Analyzing current day designs to benchmark the state of the art and to identify processor subsystems that are potential power hogs.
3. Development of new techniques that will reduce these metrics.
4. Demonstration of these techniques in a prototype CMOS PISA CPU chip.

These new techniques will include gate design level approaches, but will focus primarily on organizational and instruction set architectural approaches that inherently have a lower power requirement. The end goal of this research is to develop a deeper scientific understanding of the relationship between power and computation, and develop techniques that minimize the ratio of the two in ways that can broadly impact the continuing evolution of VLSI technology and its successful use in computing.

Iowa State University; Sachin Sapatnekar; CAREER: Gate-Level Performance Optimization of Pipelines and General Sequential Circuits; (MIP-9502556, A001); \$147,500; 36 months.

This project is concerned with performance optimization of high-speed synchronous digital circuits. It addresses the application of timing optimizations to pipelines and general sequential circuits in high-performance systems. The core of the research effort is directed towards the method of retiming at the gate level. Retiming is a sequential logic optimization technique that repositions memory elements, namely, edge-triggered flip-flops and level-triggered latches, within a circuit to optimize the timing behavior of the circuit. The use of retiming optimization methods, chiefly device sizing, is also being investigated.

Under the educational part of this project, the PI is incorporating CAD techniques in the undergraduate program from the sophomore level by the use of CAD tools in teaching. At the graduate level, he is involved in developing courses, the supporting existing courses, and in advising graduate students. The investigator is also exposing undergraduates to his research for a better appreciation of the state of the art, and to motivate them towards higher education. The PI is also contributing towards outreach activities by focusing towards distance education through teaching courses to sites in the state of Iowa via television, and around the country through the medium of satellite.

Johns Hopkins University; Smaragda Konstantinidou; CAREER: Application-Driven, High-Performance Communication in Parallel Computing Systems; (MIP-9501768); \$135,000; 36 months.

The goal of this project is two-fold. First, it seeks to study the communication characteristics of applications that can benefit from the computing power, very large memory, and high-speed communication of parallel computing systems. Within this context measurements are being taken of the traffic generated by both regular and irregular scientific applications, applications written in a variety of programming paradigms, data-mining applications and applications with real-time traffic demands. Second, the project seeks to use the information gathered from applications to guide the designs for interconnect devices and programmable communication coprocessors that efficiently support multiple classes of applications and multiple programming paradigms.

The education plan addresses two important issues in experimental computer science: First it addresses the issue to teaching a subject area, such as high performance computing, that is technology-driven and where the relevance of knowledge and information can have a very short life span. Second, it addresses the issue of teaching the methodology of experimentation and in particular the process of designing experiments.

Massachusetts Institute of Technology; Anant Agarwal; Protection and Translation in Multimodel Multiprocessors: The MIP FUGU Workstation; (MIP-9504399); \$236,265; 24 months; (Joint support with the Experimental Systems Program - Total Grant \$800,000).

This project is extending the Alewife multiprocessor to explore two aspects of future scalable multiprocessors. One aspect is protection. Protection is required on accesses to hardware resources, with as little software overhead as possible. For example, when a message sender and receiver are part of the same application, they should be able to directly access network hardware. The other aspect is address translation, so that each application can have its own address space. All processors executing an application must have the same address space, however. To ensure this, memory mapping hardware is placed at each node of the multiprocessor, and must be kept coherent.

The FUGU system provides three modes of interprocessor communication. Shared memory is provided by hardware synthesized messages among nodes, which maintain cache coherence. Short messages can be sent directly by performing loads and stores to I/O registers that are in each node's address space. Longer messages are managed by a DMA engine that is separate from the processor.

Translation and protection for shared memory and for short messages are provided by hardware enhancements that maintain coherence in the translation lookaside buffers and ensure that messages are received only by the processes that they are addressed to. For long messages, only minimal support for address translation is provided by the hardware since the need for such translation is likely to be rare. Any needed translation will be provided by software.

University of Michigan; John P Hayes; *Functional Testing of Complex Digital Systems*; (MIP-9503463); \$72,815; 36 months; (Joint support with the Design, Tools and Test Program - Total Grant \$202,000).

This research is on techniques for testing VLSI circuits directly at functional levels, while ensuring that standard testing quality goals are met. A high-level fault model called the induced functional fault (IF) model is being investigated. With this model, an IF fault class can be coupled to a physical fault class so that detection of IF faults ensures detection of physical faults. Advantages of this approach are easy derivation of near-minimal test sets, and reuse of tests in other circuits. The construction, use and reuse of precomputed tests for sequential circuits are being explored. In this regard, the application of balance and functional properties to simplifying test generation and improving fault coverage are being studied. Built-in self-testing methods that facilitate the implementation of functional testing methods and the detection and elimination of design errors are also being explored.

University of Minnesota; David J Lilja; *New Mechanisms for Parallel Loop Scheduling*; (MIP-9221900 A001); \$5,000.

Automated assignment of separate loop iterations to different processors is a powerful technique for scheduling programs to run on parallel computers. This project is investigating four research approaches to loop scheduling: investigation of processor allocation strategies, minimization of time in loop-control critical sections, techniques for scheduling loops with large but regular execution time variations, and adaptive techniques for minimizing memory reference effects. Research methods include simulation and measurement of prototype implementations on several parallel computers.

Princeton University; Niraj K Jha; *Fault Tolerance in Distributed Systems*; (MIP-9423574); \$201,893; 36 months.

This project is concerned with a novel, cost-effective scheme, called task-based fault tolerance (TBFT), for providing fault tolerance to heterogeneous distributed computing systems. The technique considers computations at the task level and places assertion checks on the tasks to achieve fault tolerance. It requires no hardware overhead and relatively small amount of delay overhead. Allocation and scheduling of the task is performed in the task graph to support fault tolerance, while at the same time an attempt is made to minimize the schedule length and obtain a balanced load. Simulation results show that the overhead in schedule length incurred due to the introduction of assertion checks is very small compared with the case when all the tasks in the system are duplicated. The topics under investigation include the following:

1. applying TBFT techniques to multiple faults,
2. developing mixed allocation/scheduling methods to reduce schedule length and fault tolerance overhead,
3. developing methods to ensure safety of operation without much adverse effect on the reliability of the system,
4. developing allocation and scheduling methods which would permit the use of redundancy at the task level to provide fast on-line diagnosis,
5. develop probabilistic diagnosis methods,
6. integrating detection/diagnosis results with a backward error recovery method to minimize the rollback domino effect using asynchronous checkpointing, and
7. developing allocation and scheduling methods for periodic and aperiodic tasks in fault-tolerant real-time distributed systems based on the notion of forward error recovery for safety-critical tasks.

Princeton University; Wayne H Wolf; *Architectural Co-Synthesis for High-Performance Distributed Embedded Systems*; (MIP-9424410); \$66,908; 36 months; (Joint support with the Design, Tools and Test Program - Total Grant \$192,383).

Embedded computing systems must be designed to meet hard, soft, performance, cost and other constraints. This research studies the synthesis of embedded systems built from multiple processors, which may be commercial or application-specific ICs connected in a network. This methodology, co-synthesis, aims to simultaneously design the hardware and software architectures of a system. Algorithms for meeting all constraints through iteration of both the hardware architecture of the distributed computing engine and the

process architecture of the application software are being investigated. The model for the research is the task graph and distributed system graph. These are being extended based on simultaneous design the hardware and software, stronger assumptions about processes in the task graph, and simultaneous scheduling, allocation, and partitioning of processes during co-synthesis. Mathematical optimization methods, such as graph partitioning and network flow, are being used to generate effective design algorithms.

State University of New York - Binghamton; Kanad Ghose; *Inherently Low Power Computer*; (MIP-9504767); \$170,001; 36 months.

This project focuses on the problem of power dissipation, primarily on the fundamental issues of what are the real sources of power dissipation in CMOS microprocessors, and what can be done to minimize it at a more global level. The goal is development and demonstration of techniques that would support Power-efficient Instruction Set Architectures (PISA). For this, there are four major tasks:

1. Developing technology independent models and metrics for power dissipation in CMOS logic.
2. Analyzing current day designs to benchmark the state of the art and to identify processor subsystems that are potential power hogs.
3. Development of new techniques that will reduce these metrics.
4. Demonstration of these techniques in a prototype CMOS PISA CPU chip.

These new techniques will include gate design level approaches, but will focus primarily on organizational and instruction set architectural approaches that inherently have a lower power requirement. The end goals of this research is to develop a deeper scientific understanding of the relationship between power and computation, and develop techniques that minimize the ratio of the two ways that can broadly impact the continuing evolution of VLSI Technology and its successful use in computing.

Ohio State University; Dhabaleswar K Panda; *CAREER: Communication and Architectural Supports for Implementing Distributed Shared Memory on Wormhole Networks*; (MIP-9502294); \$105,998; 36 months.

The goals of this project are:

1. to determine the communication, synchronization, and architectural requirements for supporting DSM; and
2. to develop design guidelines to build efficient wormhole-routed networks for supporting the DSM models.

The research is focused on four research

directions: developing multicast algorithms for effective cache-coherency, mechanisms and algorithms for barrier synchronization, techniques to alleviate hop-spot problems, and determining the necessary architectural supports. Theoretical and probabilistic analysis together with simulation experiments are being used to develop and validate the results.

The educational activities of this project include enhancing and strengthening the computer architecture curriculum (undergraduate, graduate core, and graduate research courses), formulating a capstone design course, undergraduate mentoring, and creating challenging and interactive environment for graduate learning.

Oregon State University; Shih-Lien Lu; *RIA: Implementation and Synthesis of Micropipelines in CMOS Differential Logic*; (MIP-9211510 A001); \$9,950.

The discipline of timing management is the key to the effective design of any large scale computer systems in this era of advanced technology. It has been pointed out by Sutherland recently that the current design framework which employs a globally clocked timing discipline has reached its limit. It was demonstrated that there are many advantages if a different timing discipline, namely, the transition-signal conceptual framework is used to design complex computation systems. In this research, we take a bottom up approach to accomplish the synthesis of micropipeline processing systems from a high-level specification. The first milestone is to examine the design and building of micropipeline structures using differential CMOS logic instead of static CMOS logic. The second milestone is to implement an automatic synthesis system for a particular micropipeline stage's logic processing block given a set of logic equations. The third milestone is to implement all necessary token control circuitry. The last milestone is to map a dataflow graph which specifies a particular algorithm into hardware using the micropipeline design frame work.

Carnegie-Mellon University; John P Shen; *Analytical Model of Superscalar Processor Performance*; (CCR-9423272); \$119,999; 36 months; (Joint support with the Computer Systems Architecture Program - Total Grant \$239,999).

Trace-driven simulators (TDS) are widely used by processor architects to accurately assess the expected performance of a processor design. While necessary and useful, TDS have some shortcomings. Traces of realistic programs are usually very large and expensive to store, and simulation of traces can very time consuming. Furthermore, simulation does not yield a characterization of the factors that determine performance, and hence can not recommend specific changes to enhance performance.

The two-fold objectives of this research are:

1. to investigate alternatives to TDS for accurately quantifying processor performance; and
2. to seek to characterize factors that determine superscalar processor performance.

The effort will involve developing an analytical model of machine behavior, based on which processor performance can be computed

without requiring tedious trace-driven simulation. It is hoped that the same analytical model, in addition to providing a useful tool, will also provide a characterization of superscalar processor performance. This characterization can potentially reveal fundamental laws governing processor performance that can be mathematically formulated.

In this research project, processor performance is modeled as a composition of program parallelism and machine parallelism. Program parallelism is represented by two distribution functions that are derived from the execution trace of benchmark and captures the control and data dependencies in the program. This derivation is analogous to performing a type of transform on the execution trace from the time domain to the transformed domain. Consequently the two distribution functions can be viewed as concise spectral representation of the program parallelism. Machine parallelism represents the resources available in the processor for the concurrent processing of instructions. A machine is modelled as a network, or circuit, of primitive components, e.g., buffers and scalar pipelines. Macro models of the components and the model for their serial and parallel interconnections will be developed. With such a circuit of model microarchitectures, complex superscalar processors can be effectively and rigorously modeled.

While the research appears somewhat theoretical in nature, a great deal of experimentation is required and expected. The plan is to work with realistic programs with extremely long traces as well as a wide range of state-of-the-art implementations of real instruction set architectures. The Alpha 21164 and the Power PC 620 designs and the SPEC92 benchmark suites are used as the primary experimentation vehicles.

University of Pittsburgh; Steven P Levitan; *Temporal Specification Verification*; (MIP-9102721 A003); \$25,000; (Joint support with the Design, Tools and Test Program - Total Grant \$25,000).

This research is on verifying timing specifications for interconnection of modules in both synchronous and asynchronous digital systems. The notion of temporal behavior is being abstracted from the notion of functional behavior by focusing primarily on the control protocols of the modules and ignoring the data values computed by the modules. In this model, the interface protocols of each module are given along with the connectivity between modules. A static graph is built that describes the temporal relationships among all the external signals of all the modules. The verification process is based on a comparison between the possible behaviors of the system, represented by the graph, and the legal behaviors as represented by a set of constraints. The key constraint is that the temporal behavior of one module cannot violate the temporal

constraints of another module within the system. The algorithms support multiple system states, state transitions, and checking of conditionals and loops within the protocols. This searching is tractable because functional behavior and data values generated by modules are not considered.

University of Rhode Island; Qing Yang; *Exploring the Design Space for High Performance and Low Cost Memory Hierarchies*; (MIP-9505601); \$159,818; 36 months.

The main objective of this research is to investigate the issues related to hardware designs and performance evaluations of a number of innovative cache design techniques. These design techniques are currently being developed for single-chip general purpose processors and multiprocessors. Specifically, the focus of the research includes:

1. minimization of area cost of on-chip caches by CAT--caching address tags;
2. maximization of cache hit ratios by evenly distributing data across cache sets with the help of a few tag bits that change frequently during program executions;
3. investigation of potential impacts of the new CAT cache designs for various cache configurations as well as multiprocessor caches; and
4. devising analytical models, and performing trace-driven simulations and execution-driven simulations for evaluating implementation costs, performances, and design trade-offs of various designs.

Texas A & M University; Laxmi N Bhuyan; *Cache Architectures for Large Shared Memory Multiprocessors*; (MIP-9301959 A002); \$102,384; 12 months.

This project addresses the problem of designing coherent caches for scalable multiprocessors with shared address spaces. Maintaining cache coherence in large-scale systems is time consuming due to the lack of sufficient broadcasting capacity in the interconnection networks. A dynamic cache coherence protocol is being designed and evaluated to reduce the cache coherence overheads in large systems. This protocol limits invalidation or update traffic to a subtree of the interconnection network. The first phase of the project consists of embedding rooted trees into various networks, and studying the hardware and timing complexities of directories using the embeddings. In a second phase, issues such as fault tolerance adaptive routing, and task mapping are being addressed. The research makes use of analytic techniques and execution driven simulation.

Texas A & M University; Nitin Vaidya; *CAREER: Two-Level Failure Recovery Schemes for Multicomputers and Distributed Systems*; (MIP-9502563); \$85,978; 36 months.

The thrust of this proposal is on backward recovery schemes for multicomputer and distributed systems. The research is based on a two-level approach that uses a low overhead recovery scheme to tolerate more probable failures, while the less probable failures may

be tolerated with a higher overhead. By minimizing the overhead for the more frequently occurring failure scenarios, the proposed two-level approach will achieve lower average performance overheads as compared to existing recovery schemes. Objectives of the research are:

1. the design of efficient two-level recovery schemes,
2. the development of optimization techniques for two-level recovery schemes to determine various parameters (e.g., checkpointing frequency) that will minimize the average performance overhead for the recovery scheme,
3. the implementation of two level recovery schemes on a 64 processor nCUBE-2 multi-computer and on a network of workstations, and
4. the experimental evaluation of two-level recovery schemes to demonstrate their ability to achieve better performance than the existing recovery schemes.

The goal of the proposed education activities is to incorporate the results of recent research into undergraduate and graduate curriculum. Specific objectives for the education plan are:

1. the development of an undergraduate introductory course on fault tolerance and a graduate course on fault tolerance techniques for parallel and distributed systems;
2. the development of an associated undergraduate research effort; and
3. the development of a software toolkit to facilitate experiments in the proposed courses.

University of Utah; Erik L Brunvand, Ganesh Gopalakrishnan; *The Design of Asynchronous Circuits and Systems with Emphasis on Correctness and Proven Optimizations*; (MIP-9215878 A001); \$42,610; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$85,220).

This research merges two efforts in asynchronous circuit compilation. These are the work of Gopalakrishnan on the language hopCP and its use in verification; and the work of Brunvand on asynchronous circuit compilation. The research is:

1. enhancing the expressive power as well as the semantic clarity of concurrent hardware description languages for asynchronous circuits and systems;
2. extending the formal basis for compiling from HDL's to circuit designs;
3. formally characterizing and improving the optimizations used in asynchronous circuit compilation; and
4. studying the performance of implemented circuits with regard to a variety of parameters.

University of Utah; Ganesh Gopalakrishnan; *A Multi-Paradigm Verification System Tailored for the Design Refinement Cycle*; (MIP-9321836 A001); \$5,000.

The design of a VLSI system involves multiple design representations; and the design must go through several iterations aimed at meeting many performance and cost constraints. Verification that the design meets constraints is necessary. This research is developing rigorous verification methods that span multiple design representations, accommodate design revisions, and provide incisive partial verification methods (e. g. verification focussed on the "corners" of the behavioral space) that fit within designers' time budgets. These ideas are being validated by verifying the real asynchronous designs.

University of Vermont; Yuanyuan Yang; *Routing Control Strategies for Multicast Networks*; (MIP-9522532); \$18,000; 12 months.

This is a Research Planning Grant for women who have not had prior independent Federal research. The award is allowing the PI to develop a research program in the area of multicast communications for highly parallel computing systems. The planning activities are concentrated on the investigation of routing control strategies for multicast networks from the following aspects:

1. designing new routing control strategies which can more effectively reduce the non-uniformity of multicast connections;
2. developing a network simulator to simulate the multicast networks under these routing control strategies, and gather a large amount of statistical performance data to compare different control strategies and guide the theoretical analysis in the later phase of the research systems.

Virginia Polytechnic Institute and State University; James R Armstrong, Walling R Cyre; *Rapid Development and Testing of Behavioral Models*; (MIP-9120620 A005); \$20,000; (Joint support with the Design, Tools and Test Program - Total Grant \$20,000).

This research is on methods to create behavior models that accurately represent the functionality and timing of complex devices. The work is on developing a "Modeler's Assistant" as a base for structured

development of behavioral models. Specific problems being solved are: developing a process primitive set for the Modeler's Assistant, developing and evaluating performance measures for structured behavioral model development, developing a natural language interface for the Modeler's Assistant, and building into the Modeler's Assistant the capability to automatically generate tests for any behavioral model which has been constructed by the system. Behavioral models are being expressed in the high-level language, VHDL.

Application Driven Architecture

University of Arizona; Ahmed Louri; *Development and Realization of an Optical Content Addressable Parallel Processor for High-Speed Database Processing*; (MIP-9505872); \$335,152; 36 months.

This research is exploring the application of optical systems for high-speed database processing. It is combining a parallel model of computation, associative processing based on the use of content addressable memory, with an inherently parallel technology, optics. The purpose is to integrate into one study the elements of parallel architectures, parallel algorithms, device technology, and laboratory demonstrations which will test and demonstrate the advantages of optics for high-speed database processing.

The approach taken on this project provides direct hardware support for database operations. It consists of identifying the basic fundamental operations required for data base processing and parallelizing them as much as possible. This is followed by developing an optical parallel architecture that directly implements the parallelized operations. Parallel algorithms for database processing are then developed that build on these fundamental operations. Consequently, the optical architecture will support database applications in a truly parallel fashion.

California Institute of Technology; Michelle Effros; *CAREER: Code Clustering for Universal Image Coding and Other Implications*; (MIP-9501977); \$32,169; 36 months; (Joint support with the Systems Prototyping and Fabrications Program, the Circuits and Signal Processing Program, the Design, Tools and Test Program - Total Grant \$134,628).

This project seeks to develop reasonable complexity, source-independent coding algorithms, which are crucial to the design of robust systems for image coding and mobile communications. In these applications the statistics of the source and channel in operation are typically unknown a priori, and the

performance of the coding strategy employed is sensitive to those unknown characteristics.

The two-stage approach developed in the source coding literature is employed. The literature demonstrates that in general one should quantize the space of possible codes. Some of the rate should be spent on describing which code, in a family of codes, should be used on the source in operation. Specific projects include the development of a universal DCT code compatible with JPEG and MPEG image and video standards, a universal KLT code, a universal wavelet packet code, and a universal channel code.

The main objectives of the education plan are to develop and maintain an exciting atmosphere for active learning for undergraduate and graduate students through innovative programs that encourage the maximum possible exchange between students, faculty, and individuals from local industry.

San Jose State University; Belle Wei; *RUI: A VLSI Arithmetic Data Path for a Complex Number Digital Signal Processor*; (MIP-9321143 A001); \$35,285; 12 months.

This project deals with the design and implementation of VLSI algorithms and circuits for arithmetic operations on complex numbers. The research is motivated by the need for these circuits in high-speed digital signal processing applications. The goals of the research are to:

1. Develop VLSI design algorithms for the arithmetic data path of complex number digital signal processors.
2. Design and implement VLSI circuits and components for the developed algorithms, with emphasis on both architecture and circuit design.
3. Study the trade-off between chip area and VLSI circuit latency/throughput and to develop a design curve exhibiting their trade-offs for benchmark complex number applications.
4. Specify other data path components for the complex number digital signal processor to support high-throughput input/output operations and ease of programming.

University of California-Irvine; Kai-Yeung Siu; *NYI: Analysis and Design of Artificial Neural Networks*; (MIP-9357553 A002);

\$31,250; 12 months; (Joint support with the Circuits and Signal Processing Program - Total Grant \$62,500).

Artificial neural networks present a new model for massively parallel computation and a promising paradigm for solving large scale optimization problems. This research is exploring the advantages of neural network-based models over conventional models for computation, and a novel design of neuromorphic computing architectures for applications in signal and image processing. A theoretical framework is being established to derive tight tradeoffs between the number of elements and the number of layers in neural networks. The results should answer some of the key open questions in the analysis of neural networks using classical mathematical tools such as rational approximation techniques and harmonic analysis.

University of South Florida; N. Ranganathan; *VLSI Architectures for Pattern Matching and Recognition*; (MIP-9407034 A001); \$28,220.

The goal of this research is the investigation, design and implementation of high performance VLSI architectures for two dimensional pattern matching and recognition. Although several hardware algorithms have been proposed in the literature for the one-dimensional pattern (string) matching problem and its variations, very little work has been done in the area of hardware algorithms for matching and recognition of two-dimensional patterns. Most of the research has been aimed at efficient software algorithms based on different features and matching strategies. Many pattern matching applications need real-time response, perform number of repetitive computations on different data sets, make use of regular and local operations, are modular, and have a fair amount of inherent parallelism. The design of special purpose hardware for pattern matching could speed up the task considerably, making it amenable for real-time applications.

Specifically, the main objectives of this research are to develop new and efficient hardware algorithms and VLSI chips for a class of pattern matching problems such as scene matching, approximate string matching, polygon matching and tree matching.

Design and develop an application specific system architecture for pattern matching and recognition through integration of the above mentioned VLSI components at board level.

Illinois Institute of Technology; Wai-Yip G Chan; *CAREER: Audio-Visual Signal Compression Using Vector Quantization*;

(MIP-9502629); \$52,520; 36 months; (Joint support with the Circuits and Signal Processing Program - Total Grant \$112,041).

This research involves the development of efficient, reliable, and real-time realizable coding algorithms and techniques for compressing and transporting audio-visual signals, targeting particularly the visual-telephone application. For a total channel bit-rate of 64 kb/s or less, our goal is to achieve superior grade-of-service by treating the compression of the audio (speech) and video signals as an integrated problem. Our thrust centers around the use of vector quantization techniques, and schemes that jointly optimize the signal analysis and quantization functions. We place particular emphasis on finding an efficient representation of the displacement ("motion") field for motion compensated prediction of the video signal. To facilitate personal wireless communications, we are exploring adaptive coding schemes that exploit the prioritized transmission of encoded data. For compressing speech signals, we are exploring variable rate coding schemes, and are allocating bits between the speech and video signals to optimize the overall grade-of-service. Achieving our goal would give significant impetus to the development of the emergent Information Superhighway, which can form part of an infrastructure for supporting high performance computing.

Educational activities undertaken in conjunction with the above research include but are not exclusive to: the development of undergraduate digital signal processing laboratory courses; the introduction of graduate courses in the area of audio-visual information processing and communications; increasing the use of computer-based multimedia instruction in the classroom; assisting in speeding up the digitization of the library and networking of its databases.

Johns Hopkins University; Andreas G Andreou, Fernando Pineda; *Analog Computation and VLSI Architectures for Contraction Mappings*; (ECS-9313934 A002); \$37,878; 12 months; (Joint support with the Computational Engineering Program - Total Grant \$75,756).

This project will attempt to develop a new class of recurrent networks. The architecture of the networks is inspired by recent work on image encoding based on iterated transformation theory and its associated inverse problems. The PIs purpose to restrict our investigation to networks that can be physically implemented in subthreshold analog VLSI. With their approach analog components that implement high quality arithmetic operations are unnecessary. Indeed, significant departures from ideal linear behavior can be tolerated, provided that these departures are reproducible across chips. As a concrete application they will consider that task of data compression and decompression. Compression is accomplished by the relaxation of an electronic circuit to a steady state while compression is preformed either off-line or with an adaptive analog VLSI neural network architecture. For hardware compression they propose to use a learning algorithm that learns both the weight and the connection topology. Hence, the ability to gate and switch electrical current is central to the operation of these networks. The main thrust of this investigation is to design and characterize the circuits that implement the required transformations in one dimension. A successful outcome to this investigation has the potential for making compression and decompression technology available for all low-power applications.

Northeastern University; David R Kaeli; *CAREER: Architectural Support for Object-oriented Code Execution*; (MIP-9501172); \$131,995; 36 months.

This research focuses on the inherent performance problems associated with object-oriented code execution. Traces are being captured from a number of object-oriented programming environments. A workload characterization is being performed to identify the underlying problems presented by this new programming model. Architectural changes are then proposed to improve the execution performance of object-oriented code. This research is investigating architectural changes to the hardware that will improve performance across all programming environments. As more and more software is developed using object-oriented languages, this research will have a direct impact on acceleration the performance of future applications. The educational part of this project is concerned with making improvements in the Computer Engineering curriculum, student advising tools, classroom computer aids, and undergraduate research

courses.

New Jersey Institute of Technology; Zoi-Heleni Michalopoulou; *Signal Processing for Marine Mammal Localization and Deconvolution of Biological Acoustic Transients*; (MIP-9505362); \$14,858; 12 months.

Concerns about the effects of human interference in the oceanic environment upon the life and habits of marine mammals have been recently raised. In response, researchers have taken interest in processing all relevant information to better understand this interference with the hope and goal of protecting marine life. In this context, the development of signal processing methods for the exploration of sounds generated by marine mammals in the ocean is being pursued. Success of this research will enable the tracking of mammals, and the reconstruction of the signals they produce, through remote acoustic sensing.

Localizing marine mammals in the ocean and identifying the sounds they generate requires the establishment of a link between physics (acoustics) and signal processing. The normal modes approach is being adopted for modeling the acoustic signal propagation in the oceanic environment. Models will then be sought that capture the time and frequency structure of biological signals. The combination of acoustics and signal analysis is expected to facilitate inversion of the acoustic field measured at an array in the ocean for the estimation of the location of sound-transmitting marine mammals and the actual form of the signals they transmit.

Rutgers University New Brunswick; Roy D Yates, Christopher Rose; *Power Control For Packet Radio Networks*; (NCR-9506505); \$37,000; 12 months; (Joint support with the Communications Research Program, the NSFNET Program, and the New Technologies Program - Total Grant \$174,106).

This research looks to extend the application of power control in mobile cellular communications to multihop CDMA packet radio networks. Using soft interference constraints derived from signal to interference ratio measurements, distributed power control algorithms will be studied which adapt to changes in the radio link quality, actively regulate the network topology, achieve desired end-to-end throughput objectives, and integrate power control and multihop routing.

Oregon State University; Bella Bose; *Balanced Codes for VLSI Systems*; (MIP-9404924 A001); \$5,000.

This project deals with the properties and applications of balanced codes. In a balanced code, each code word contains equal number of 1's and 0's. These codes find many applications in computer and communication systems such as noise reduction in VLSI systems, fault masking in bus lines of VLSI systems, delay insensitive communications in asynchronous systems, data transmission in fiber optics, data storage in optical discs and magnetic tapes, and fault tolerant synchronous circuits. The objective of this research is to develop design methods for balanced codes suitable to

these applications. Our aim is to design codes which require low redundancy but at the same time have fast and simple encoding/decoding algorithms. Other topics such as error correcting balanced codes, asymmetric error correcting and detecting codes applicable to the above mentioned applications will also be investigated.

Brown University; Harvey F Silverman; *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A002, A003, A004); \$200,000; 24 months; (Joint support with the Systems Prototyping and Fabrications Program, the Experimental Systems Program, the Circuits and Signal Processing Program, and the Design, Tools and Test Program - Total Grant \$674,601).

This is a joint project between Brown and Rutgers to build a large microphone array with associated processors for beam forming. Applications are direction finding, echo cancellation, and speaker differentiation in teleconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

Brown University; Harvey F Silverman; *Parallel Architectures for Speech Recognition: Nonlinear Optimization of Expectation-Maximization (EM) Training of Hidden Markov Models (HMMs) in a Reconfigurable Environment*; (MIP-9509505); \$136,436; 12 months; (Joint support with the Circuits and Signal Processing Program - Total Grant \$187,251).

This research focuses on one of the strategic areas of national concern, that of high performance computing. It involves the development, construction and testing of new architectures for high-speed computing. The idea is to combine general-purpose, RISC-based processing nodes, each with multiple field-programmable-logic-array (FPLA) based coprocessor systems. This kind of system has the hardware properties of a general-purpose computer, but the advantage of performance more akin to that of a special-purpose engine. In particular, the 20-node Armstrong III system has been built, is operational, and provides nearly two-orders of magnitude improvement over current advanced workstations. The hardware system, combined with its configuration compiler, a program that automatically translates a C function subroutine to both machine code and to the hardware design of the reconfigurable coprocessors, are the basic building blocks for this class of architectures.

This project is unique in that not only has the hardware/software system been built, but it is also being tested on real applications. The training of a modern Hidden Markov Model (HMM) based speech recognition system requires hundreds of hours, even with recent variants developed at Brown and elsewhere. The reduction of this training time to ten or so minutes allows progress in this area to be made at a faster rate and/or that expensive nonlinear optimization techniques may now be applied to the problem. Also, this means that data from a less cumbersome sensor system (the microphone array systems in place and being developed at Brown) can be suitably incorporated into a more robust speech recognition system.

Texas A & M University; Nitin Vaidya; *Bit/Byte Bounded Error Control Codes for Byte-Organized Systems*; (MIP-9423735); \$47,701; 12 months.

The principal goal of this project is to design error control codes for "byte-organized" systems using a new error control model. In a byte-organized system, the hardware is partitioned such that each part produced a subset of hits in the output. For example, a memory may be organized as one byte per card. In this case, each card produces a byte in the memory word. For such systems, a new error control model called the bit/byte bounded model is being developed. Briefly, this model can be described using two parameters, say t and u . The codes designed using the bit/byte bounded model can tolerate up to t bit errors confined to a must u bytes. Unlike the existing approaches, the bit/byte bounded model can interpolate between the traditional bit and 1 byte error models. Also, the proposed model can be used to reduce the number of checkbits by trading a small amount of error control capability.

University of Wisconsin; Parameswaran Ramanathan; *Time-Constrained Communication in Real-Time Systems With Point-to-Point Interconnection Topology*; (MIP-9213716 A001); \$5,000.

A major limitation of real-time distributed computing systems is that communication latency between tasks on different nodes can be large. This research is directed towards alleviating this limitation using three methods. First, algorithms for assigning deadlines to

information transfers between tasks are being developed. Second, partitions of information transfers into messages are being investigated. Third, routing and delivery strategies are being developed. The result is a set of solutions which can be used to design real- time communications systems.

Workshops and Conferences

University of Washington; Jean-Loup Baer; *Travel Support for ISCA '95, June 22-24, 1995, Santa Margherita Ligure, Italy; (MIP-9521566); \$15,000; 6 months.*

This is an attendance and travel grant for the 22nd Annual International Symposium on Computer Architecture on June 22-24, 1995 in Santa Margherita Ligure, Italy. It is co-sponsored by the Association for Computing Machinery and the Institute of Electrical and Electronic Engineers. The symposium, through a combination of invited talks, panel sessions, tutorials, workshops and refereed paper presentations, continues to serve the various needs of the computer architecture research community. It is an important conference in the computer architecture area. This grant provides support to help several graduate students attend the symposium.

Circuits and Signal Processing

Dr. John H. Cozzens, Program Director
(703) 306-1936 jcozzens@nsf.gov

The Program

The Circuits and Signal Processing (CSP) program supports basic research in the areas of digital signal processing, analog signal processing, and supporting hardware and software systems. This research is typically driven by important applications and emerging technologies. Signal processing is a highly active research area, with topics ranging from theory to Very Large Scale Integrated (VLSI) circuit implementations and applications. Although signal processing is typically driven by advances in technology, discoveries in this field also serve as a catalyst for new technological innovations. A taxonomy of research areas, based on signal characteristics, applications, and/or technology, include:

One-Dimensional Digital Signal Processing (1-D DSP) - the representation of time-varying signals (e.g., audio, EKG, etc.) in digital form, and the processing of such signals;

- (adaptive) filtering and equalization
- filter design, theory, and analysis, both linear and nonlinear multirate processing and wavelets
- time-frequency representations

Statistical Signal and Array Processing (SSAP) - the use of statistical techniques for the processing of signals that may arise from multiple sources;

- cyclostationary signal processing
- higher order statistics
- (statistical) array processing
- nonstationary and time-frequency

Image and Multi-Dimensional Digital Signal Processing (IMDSP) - the acquisition, manipulation, and display of multidimensional data using digital technology;

- image analysis, filtering, restoration, and enhancement image and video coding
- vector quantization

Analog Signal Processing (ASP) - the processing of data without conversion to sampled-digital form;

- analog-to-digital conversion
- analog circuits and filters

Special attention is currently given to research in:

- antenna array processing with application to wireless communications systems, especially cellular telephony, Personal Communications Systems (PCS), and wireless local area networks
- computed tomography and SAR
- data quality validation¹
- manufacturing applications, e.g., nondestructive test and evaluation
- scalable/progressive/multiresolution approaches in signal decomposition, compression, and other signal processing¹
- signal compression for reduced data rate with applications to wireless communications systems
- signal processing techniques to support content analysis¹

Low priority areas include high-level image processing and high-level speech processing, as these areas are supported elsewhere in the Foundation; and mature areas such as classical circuit theory and (classical) spectral estimation.

Participation in signal processing research by undergraduate students and underrepresented faculty sectors is encouraged. ent workshop on signal processing for the National Information

¹ This is research topic was suggested by the panelists of a rec Infrastructure (NII), sponsored by the Circuits and Signal Processing Program. This workshop was held at the National Science Foundation in Ballston, VA, on 18-19 August 1994, "to assess the current state of signal processing in the NII, to identify important issues and trends of research and development of signal processing theory and applications for the NII, and to make recommendations for research and development." The entire document - NSF'95-10 (new) - is also available at web location <http://www.ee.gatech.edu/users/215/nii/niireport.html>

Awards

Analog (Mixed Analog/Digital) Signal Processing

University of California-Berkeley; Paul R Gray, Robert G Meyer; *Research in High-Frequency Analog Electronic Circuits for Communication Systems*; (MIP-9412940); \$136,403; 12 months.

The field of electronic telecommunications has experienced explosive growth in the past ten years, both in terms of its commercial importance and in terms of the research effort devoted to it in academic and industrial laboratories. It appears that radio-based digital communications will play a much larger role than might have been anticipated, complementing data communications on wire pairs with fiber media which are coming into increased use for data communications backbones. As in the past, reducing the cost of electronics associated with VLSI technology will be critical.

This research is directed at exploring new ways to use silicon integrated circuit technology to improve the performance, reduce the power dissipation, and reduce the cost of components for communication systems of various kinds. The primary emphasis of this research is on data communications and Personal Communications Systems (PCS) using radio and free-space optical media. Particular emphasis is placed on circuit techniques applicable to implementation of RF, IF and baseband mixed signal communications circuits with low power and low operating voltage. Specific topics for investigation are: limits-to-phase noise performance in power-optimized monolithic voltage-controlled oscillators; investigation of new parallel-architecture sampling demodulators; limits-to-distortion performance in MOS sampling demodulators; the use of passive-sampling FIR switched capacitor filters for low-power, high-frequency filtering and CDMA despreading; optimization of the sensitivity in fiber optic communication receivers; the application of BiCMOS technology for the particular needs of high-speed communications; and modeling of active and passive IC components for RF design.

University of Florida; John G Harris; *CAREER: Analog VLSI Sensory Processing*; (MIP-9502307); \$107,887; 36 months.

This research deals with the implementation of continuous-time analog VLSI chips for processing of sensory signals. Two specific analog sensors are being developed:

1. a single-chip visual time-to-contact sensor that reports the amount of time until the chip collides with an oncoming obstacle (assuming constant

relative motion); and,

2. a sound localization chip that computes the direction (in two-dimensions) of a sound source using signals from two microphones.

Though these projects may appear to be unrelated, the proposed solutions using adaptive analog computation in CMOS hardware lead to many common problems that must be solved. By providing a cross-fertilization between auditory and visual processing, it is hoped that the underlying biological computational primitives and neuronal representations will be discovered. The resulting low-power analog devices have obvious applications in such areas as collision warning sensors for intelligent automobiles and improved, low-power hearing aids.

The educational plan at both the undergraduate and graduate level has a strong interdisciplinary emphasis. Too often students are taught to specialize in a narrow area and are consequently unable to relate their work to an overall system objective. Teaching objectives are:

1. to stress the interdisciplinary nature of real- world systems; and,
2. to introduce students to problem solving techniques for real-world signal-processing applications.

In existing classes, the use of laboratories and computer experiments where students can do hands-on experimentation with concepts is emphasized. New courses at the undergraduate and graduate levels that give students more exposure to real-world experimentation in laboratories will be designed.

University of Hawaii Manoa; Gregory T Uehara; *CAREER: Research and Education in Integrated Circuit Design for Communication and Magnetic Storage Systems*; (MIP-9501726); \$63,890; 36 months; (Joint support with the Solid State and Microstructures Program - Total Grant \$127,779).

There is an intimate relationship between the algorithms used in communication and magnetic storage systems, the Integrated Circuits (ICs) which implement the algorithms, and the system performance which results. As systems increase in complexity, trade-offs and innovations need to be made at all levels of a system design in order to drive system/circuit performance to approach fundamental limits. This research is pursuing teaching and research plans which have the goal of developing students into engineers prepared to make significant and important technical contributions to the development of high performance communication and magnetic storage systems.

On the side of teaching, the goal is to develop a curriculum which encourages students to make connections between both the theoretical and practical aspects of electrical engineering. The curriculum needs formal methods to bring out the synergy that exists between these two aspects of their field. On the side of research, the fundamental themes are:

1. to examine algorithms and architectures important in

communications applications;

2. to identify performance or speed limitations of conventional implementations of the algorithms and architectures from an IC implementation perspective;
3. to develop ways in which modifications may result in implementation advantages; and
4. to develop prototypes of new approaches in order to demonstrate its potential advantages.

Projects include the development of analog, digital, and mixed-signal ICs for wireless satellite and infra-red, wired LAN, and magnetic disk read channel applications.

University of Illinois; Bang-Sup Song; *Generalized Background Digital Calibration of ADC's*; (MIP-9312671 A001); \$50,422; 12 months.

This research is part of a continued effort to extend a background calibration principle to common multi-step/pipelined Analog-to-Digital Converters (ADC's). The basic idea is to replace a component trimming procedure usually done in the factory by a hidden electronic real-time calibration circuit running in background. Unlike other calibration techniques running in the foreground, this technique is based on dithering and nonlinear interpolation. The goal is to improve the performance of inherently fast ADC's by maintaining simple system architectures that perform a sophisticated trimming operation in the background. The generic research on lower-power design using a recycled residue amplifier and capacitive reference divider will help develop a family of high-performance analog digital interface circuits with low power and premium speeds not readily available in monolithic forms.

Harvard University; Woodward Yang; *NYI: VLSI Design for High Performance Signal Processing and Computation*; (MIP-9257964 A002, A003); \$125,000; 24 months.

The focus of this research is on the development of innovative VLSI design methodologies that will facilitate the next generation of high performance signal processing and computing systems. A variety of analog, digital, and mixed signal circuitry will be implemented for use in high performance computing applications including object recognition, smart sensors, adaptive neural networks, and programmable analog filters.

Oregon State University; John G Kenney; *RIA: High Performance Disk Drive Channels*; (MIP-9410172); \$90,000.

This research is examining techniques for reducing the hardware complexity for Decision Feedback Equalization (DFE) and Multi-level Decision Feedback Equalization (MDFE) in hard disk drives. Both digital and analog circuit approaches are being investigated. The first part of this work concerns the testing of a digital realization of DFE in an 0.8 mm process. The projected operating speed of this IC is about 100 Mhz. Enhancements to the digital circuitry will give improved strategies for optimal phase detection, gain detection and dc offset detection.

The second part of this work concerns the investigation of an analog implementation of MDFE. A simple architecture for MDFE, which includes both the phase and the gain detectors, is being studied for its robustness using actual playback waveforms. From this investigation, the circuits needed for the transversal filters comprising MDFE will be obtained. A mixture of switched-capacitor and current-mode circuit techniques will be examined.

Oregon State University; Richard Schreier; *RIA: Fundamentals of Delta-Sigma Modulation*; (MIP-9210935 A003); \$30,000; 12 months.

Delta-sigma modulation forms the foundation for the highly linear and manufacturable analog-to-digital and digital-to-analog converters known as "oversampled noise-shaping" converters. These circuits find application in narrow band systems such as digital audio equipment and medical and geophysical instrumentation. At present, there are no adequate tests for stability of these nonlinear systems - designers can only use extensive computer simulations and hope that the simulations exercise the circuit adequately. This research is developing a quick and completely rigorous method which combines analytical techniques with numerical algorithms to prove the robust stability of a delta-sigma modulator. The idle-channel noise of delta-sigma modulators is a second research topic. How to guarantee aperiodic behavior in delta-sigma modulators is being shown, but this alone is not sufficient to guarantee an absence of tones. Listening experiments using data derived from simulations need to be performed in order to judge the effectiveness of the technique. The results of these investigations are being incorporated into a computer program for the automated design of delta-sigma modulators.

Washington State University; Terri S Fiez; *NYI: High Performance Analog Signal Processing for Mixed-Mode*

The focus of this research is the development of high performance analog architectures and circuits for mixed analog-digital IC's. Emphasis will be placed on improving the immunity of the analog circuits to digital switching noise and in obtaining high speed and accuracy analog circuit performance with the newly emerging 3.3 Volt power supply standard. Current-mode circuits will be examined as a way of overcoming these limitations and this research will yield a general understanding of the advantages and limitations of current-mode circuits as compared to voltage-mode circuits. The current-mode circuits will be used to implement area efficient, high-order sigma-delta A/D converters.

One-Dimensional Digital Signal Processing

Auburn University; Jitendra K Tugnait; *Higher Order Statistical Signal and Image Processing and Analysis*; (MIP-9312559 A001, A002); \$58,875; 12 months.

This research is concerned with development, analysis and evaluation of algorithms for signal/image processing in addition to, or in lieu of, the usual second order statistics. Both time series (only system output is observed) and system identification (both input and output are observed) formulations are being considered. One-dimensional and multidimensional signals and systems are both being investigated. Whereas the Second Order Statistics of signals are a function of only the underlying system transfer function magnitude, Higher Order Statistics (HOS) of the data carry useful information about the phase characteristics of the underlying signal/system. This is crucial in deconvolution problems such as those arising in digital communication channel equalization, seismic wavelet processing, and image restoration, analysis and synthesis.

Time domain as well as frequency domain methods using higher order cumulant functions and higher order cumulant spectra (or higher order integrated cumulant spectra), respectively, are being pursued. Among the applications being investigated are:

1. image texture synthesis and classification;
2. differential time-delay and doppler estimation in unknown spatially correlated Gaussian noise;
3. blind deconvolution with unknown, possibly nonminimum phase, channels including image restoration; and,
4. system identification with noise inputs.

University of Southern California; Jerry M Mendel; *Applications of Fuzzy Systems to Signal Processing*; (MIP-9419386); \$207,254; 36 months.

For many practical problems, the information concerning the system under study is often represented in two forms: one is a set of input-output data pairs, and the other is a set of linguistic descriptions about the system (often in the form of IF-THEN fuzzy rules). Traditional system identification methods and classification methods, even including the (model-free) neural network approaches, can only utilize the input-output data pairs. This research will continue to develop general methods to combine both input-output data pairs and linguistic IF-THEN rules into signal processing systems designs.

This research is taking a new direction in handling uncertainty by using fuzzification in a fuzzy logic signal processor. The nonsingleton Fuzzy Logic Systems (FLSs) are being applied to a variety of important signal processing problems including fuzzy classification.

The objectives are to:

1. establish a measure of uncertainty for the output of a nonsingleton FLS, in accordance with estimation theory practice;
2. develop a nonsingleton FLS state estimator for nonlinear dynamical systems;
3. extend fuzzy logic classifiers to include imprecise training samples and noisy measurements, using nonsingleton FLSs;
4. extend backpropagation and orthogonal least squares parameter training procedures to nonsingleton FLSs and classifiers;
5. apply nonsingleton FLSs to the forecasting of multiple time series and to adaptive filtering problems of noise cancellation and blind equalization; and,
6. apply nonsingleton fuzzy logic classifiers to modulation classification.

Colorado State University; Richard A Davis, Peter J Brockwell, Murray Rosenblatt; *Mathematical Sciences: Time Series Models*

and Extreme Value Theory; (DMS-9504596); \$10,000; 12 months; (Joint support with the Statistics Program - Total Grant \$70,000).

The research is concerned with problems of estimation and research for time series models whose theory is not yet fully understood. The standard linear Gaussian models for time series data inadequately describe many of the time series observed in practice. It is therefore important to develop techniques for estimation and prediction based on more general models. Efficient estimation procedures and prediction techniques for non-causal and non-invertible ARMA processes are being developed and a study of the theory and application of continuous-time linear and non-linear ARMA processes is being made. Extreme value theory for linear and non-linear models is also investigated.

Existing methods of forecasting are based on assumptions which are frequently not satisfied by observed economic and scientific time series data. This research develops methods of analysis and forecasting for a more general class of time series models, leading to a more accurate forecasting of series which do not meet the restrictive assumptions of the classical theory.

University of Colorado; Delores M Etter; *Adaptive IIR Filtering Using a Stochastic Filter*; (MIP-9106126 A004); \$16,000; 12 months.

A stochastic filter consists of a bank of fixed filters with a set of corresponding probabilities. The fixed filters form a basis set of filters for the stochastic filter, and the probabilities determine the specific realization represented by the stochastic filter. This research is investigating the use of a stochastic filter to adaptively model an Infinite Impulse Response (IIR) system. Guidelines for selecting the basic set of filters are being developed in order to represent an adaptive IIR filter and to meet both accuracy and convergence speed constraints.

Georgia Institute of Technology; Petros Maragos; *Morphological Signal Processing: Slope Transforms, Max-Min Dynamics, and Differential Morphology*; (MIP-9421677); \$228,886; 36 months.

Morphological systems are a broad class of nonlinear systems that can provide rigorous and efficient solutions to many applications of image and signal processing that can benefit from and most often require the use of nonlinear analysis. Examples include multiscale filtering, feature extraction, image segmentation, and other geometry-based problems in computer vision. This research is investigating a range of theoretical and applied problems including:

1. the modeling of multiscale processing in image analysis, image distance transforms, and signal envelope detection using morphological systems;
2. the development of analytic tools for these nonlinear systems both in the time/space domain and in a new transform domain - the slope domain; and
3. the exploration of exciting connections of the above to physics.

The unifying theme is a collection of max-min differential/difference equations modeling the scale or time/space dynamics of morphological systems, and some novel nonlinear signal transforms, called slope transforms.

Georgia Institute of Technology; Petros Maragos; *Nonlinear Systems for Speech Signal Processing*; (MIP-9396301 A002); \$54,696.

This proposal focuses on the development of nonlinear signal processing systems and algorithms that can model or extract information about two types of nonlinear and time-varying phenomena in speech production: modulations and turbulence. A model is proposed for short-time speech resonances that combines both amplitude and frequency modulation. Preliminary experimental findings are consistent with this model. Fractal models for describing the geometric fragmentation of the speech signals will be used to quantify the degree of speech turbulence. A nonlinear filtering algorithm is being used to measure the short-time fractal dimension of speech signals.

Cornell University; C. Richard Johnson; *Blind Adaptive Fractionally-Spaced Linear Equalizer Behavior*; (MIP-9509011); \$221,636; 36 months.

The central theme of this research is the development of a fuller theoretical description of the behavior of memoryless-error-function (or Bussgang) type blind equalization algorithms, in particular, the Constant Modulus Algorithm (CMA), as used in practice with finite-length, fractional spacing, higher-order (non-constant modulus, non-uniformly distributed) source constellations, (deterministic and stochastic) correlated source time series, and modest channel noise. The initial target is the behavior of fractionally-spaced CMA with a correlated (or periodic) source, which appears capable of causing misbehavior in practice. In addition to averaging theory tools popular in adaptive systems analysis, it is planned to use analytical tools from topologically based, algebraic-geometry (in particular Bernstein's theorem, relaxation methods, and Morse theory) non typically associated with the study of adaptive filters. The ultimate intent is to convert an expanded theory into situation dependent design guidelines for stepsize, length, fractional spacing, initialization selection, and failure recovery tricks for the use of CMA as a start-up scheme (with a subsequent switch to decision-direction) for high data throughput applications. Analytical advances on CMA are being translated to fractionally-spaced realizations of other blind equalizer algorithms of the memoryless error function class (including refinements to CMA), for which actual operating data can be obtained industrially as for CMA. Behavioral insights into CMA are being used to help provide fair comparisons on actual operating data to competing schemes such as symbol-spaced decision feedback equalizers and similar complexity algorithms based on second-order correlation statistics of single-input, multiple-output models of fractionally-spaced equalization.

Ohio State University; Peter Clarkson; *Robust Parameter Estimation in the Detection of Cardiac Arrhythmias*; (MIP-9220769 A001); \$75,583; 12 months.

This research is concerned with improving the efficiency and robustness of parameter estimation schemes employed by implantable and portable devices for cardioversion and defibrillation. The goal of the research is to replace the ad-hoc estimation schemes used in present generation devices by robust estimators that produce optimal minimum variance estimates for a wide range of unknown error distributions, and which maintain close-to-optimal performance in the presence of parameter fluctuations. These objectives are pursued through estimation schemes that are based on the

application of Order Statistic (OS) operations to the data.

Carnegie-Mellon University; Virginia L Stonick; *PYI: Practical Approaches to Optimal Adaptive Filtering of Real Time Non-Linear Systems*; (MIP-9157221 A007); \$62,500; 13 months.

This research addresses the use of numerical optimization methods to develop real-time adaptive filters for estimating, identifying, or predicting time-varying and potentially nonlinear processes. The first phase of this work is devoted to the development, analysis, and simulation of an optimal adaptive infinite-impulse response (IIA) filtering algorithm for telecommunications using homotopy continuation methods to perform the necessary nonlinear optimization. This research will increase our understanding of IIA filter structures in time-varying environments, and will ultimately lead to their more widespread use.

William Marsh Rice University; Richard G Baraniuk; *NYI: Signal Analysis and Processing in Matched Coordinate Systems*; (MIP-9457438 A001); \$62,500; 12 months.

This research aims to extend current methods of time-frequency and time-scale analysis by developing a general theory for signal analysis and processing in alternative coordinate systems. Specific tools under investigation include optimal, signal-dependent time-scale representations, information measures for time-frequency and time-scale analysis, and operator-based, generalized coordinate systems. Test signals are being drawn from problems in machine health monitoring, magnetic resonance imaging, and dispersive signal processing.

University of Utah; Scott C Douglas; *CAREER: Advanced Architectures for Multichannel Active Noise Control*; (MIP-9501680); \$121,465; 36 months.

Active noise control is the suppression of unwanted noise via destructive interference using a digital signal processing system with sensors and actuators. The goals of this research are to achieve significant performance gains and computational complexity reductions in active noise control technology through fundamental research in multichannel algorithm design and analysis to enable the wider use of this technology in industrial and consumer markets. These efforts include:

1. the development of computationally-efficient, gradient-based adaptive controllers employing simplified updating strategies;
2. the application of least-squares and orthogonalization methods for robust and statistically-efficient filter adaptation;
3. the design of multichannel subband implementations of active noise controllers using adaptive infinite impulse-response filters; and
4. analysis and development of multichannel system identification methods that require a minimum of processing capability and that provide robust performance in time-varying situations.

The research effort includes both system simulation and hardware

development to verify the underlying principles, performances, and constraints of real-world active noise control systems.

To be most useful, an engineering education must forge a relationship between the capabilities of various technologies and the needs of the community at-large. The goal of this effort is to bring the industrial, societal, and academic communities closer together through individual and group interaction. Activities to foster this interaction include:

1. promoting joint university/industry technical projects through participation in the University of Utah Engineering Clinic Program;
2. incorporating research activities in active noise control within the undergraduate and graduate curriculum through individual class projects;
3. developing an integrated electrical engineering curriculum through departmental committee work; and
4. reaching out to the community through Engineering Career Days and the advising of prospective transfer students.

The educational effort involves the design of new course projects, committee work at the Department level, and extensive interaction with students, other educators, and members of the local community.

Image and Multidimensional Digital Signal Processing

California Institute of Technology; Michelle Effros; *CAREER: Code Clustering for Universal Image Coding and Other Implications*; (MIP-9501977); \$38,121; 36 months; (Joint support with the Systems Prototyping and Fabrications Program, Microelectronic Systems Architecture Program, Design, Tools and Test Program - Total Grant \$134,628).

This project seeks to develop reasonable complexity, source-independent coding algorithms, which are crucial to the design of robust systems for image coding and mobile communications. In these applications the statistics of the source and channel in operation are typically unknown *a priori*, and the performance of the coding strategy employed is sensitive to those unknown characteristics.

The two-stage approach developed in the source coding literature is employed. The literature demonstrates that in general one should quantize the space of possible codes. Some of the rate should be spent on describing which code, in a family of codes, should be used on the source in operation. Specific projects include the development of a universal DCT code compatible with JPEG and MPEG image and video standards, a universal KLT code, a universal wavelet packet code, and a universal channel code.

The main objectives of the education plan are to develop and maintain an exciting atmosphere for active learning for undergraduate and graduate students through innovative programs that encourage the maximum possible exchange between students, faculty, and individuals from local industry.

Stanford University; Robert M Gray, Richard A Olshen; *Tree-structured Image Compression and Classification*; (MIP-9311190 A003); \$80,054; 12 months.

Tree-structured vector quantization is an approach to image compression that applies ideas from statistical clustering algorithms and tree-structured classification and regression algorithms to produce compression codes that trade off bit rate and average distortion in a near optimal fashion. This research is examining the explicit combination of these two forms of signal processing, compression and classification, into single tree-structured algorithms that permit a trade off between traditional distortion measures, such as squared error, with measures of classification accuracy such as Bayes risk. The intent is to produce codes with implicit classification information, that is, for which the stored or communicated compressed image incorporates classification information without further signal processing. Such systems can provide direct low level classification or provide an efficient front end to more sophisticated full-frame recognition algorithms. Vector quantization algorithms for relatively large block sizes are also being developed emphasizing multiresolution compression algorithms. In order to improve the promising performance found in preliminary studies or combined compression and classification, it will be necessary to use larger block sizes or, equivalently, more context. Multiresolution or hierarchical quantizers provide a simple and effective means of accomplishing this. Other related issues are being explored, including improved prediction methods for predictive vector quantization and image sequence coding.

University of California-Berkeley; Martin Vetterli; *Adaptive Signal Decompositions with Applications in Compression*; (MIP-9321302 A001, A002); \$134,992; 12 months.

This research explores adaptive methods for high performance, lossy signal compression, and in particular, focuses on a variety of methods that rely on signal expansions. Among these are:

1. Signal adaptive expansions that include quantization and entropy coding. These are generalizations of transform or subband/wavelet and wavelet packet schemes. In particular, adaptive wavelet packets are being studied, both from the point of a view of constructing bases and that of finding good and efficient algorithms to find the best bases.
2. Overcomplete expansions or frames. The focus is on quantization performance and computational complexity.
3. The class of compression algorithms that use successive approximation. A rate-distortion version of matching pursuit is being developed; generalizations of hierarchical methods based on wavelets and wavelet packets are being investigated.
4. Adaptive schemes for compression, including adaptive transforms and quantization. Goals include a lossy equivalent of arithmetic coding, and a lossy dictionary based predictive compression scheme that resembles a lossy Lempel-Ziv algorithm.

University of Southern California; Antonio Ortega; CAREER: Adaptive Compression Techniques for Digital Video Communications; (MIP-9502227); \$135,000; 36 months.

The area of digital image and video compression and communications has seen growing activity in recent years. Activity has occurred in standardization efforts such as JPEG and MPEG, and in increasing the number of current or proposed applications from Cable or Satellite TV to video conferencing and future multimedia services. This trend will continue to create an incentive for new compression techniques that are more efficient as well as better adapted to the particular transmission environments. The main thrust of this research is to provide adaptive compression techniques for video communications.

Adaptive techniques are those where encoders are capable of matching their encoding procedure to the local characteristics of the source or the channel. Rate-distortion optimal solutions and fast heuristic approximations are being obtained for some of the problems of interest in video communications, such as bit allocation and rate control. A novel scalar adaptive quantizer, close in spirit to arithmetic coding, is being investigated. This new technique, combined with such popular quantization techniques as trellis coded quantization, will be used in real encoding situations. Convergence and asymptotic properties of this scheme are also being studied. For variable channel situations

where transmission resources are shared, as in ATM networks for instance, rate constraints for each video connection are being designed to enable good individual video quality as well as efficient overall network utilization.

This research is being carried out in the Digital Video Communications research laboratory that is currently being set up within the Signal and Image Processing Institute. The main education goal is to provide teaching in the area of signal, image and video processing. A digital video communications course, designed to cover material that is specific to video compression including the relevant communications aspects, and to provide hands-on experience for the students is being introduced.

Georgia Institute of Technology; Ronald W Schafer, S. J. McGrath, Mark A. Clements, James H. McClellan, Thomas P. Barnwell; Infrastructure and Research Program for Signal Processing in Multimedia Systems; (MIP-9205853 A005); \$168,782.

The objective of this proposal is to establish an infrastructure for signal processing in multimedia systems. Equipment will be acquired, installed, and integrated into an existing research environment consisting of networked UNIX workstations and minicomputers. This equipment will provide the capabilities to acquire multimedia signals, including image, video, speech, audio, text, and fax; store and access these signals on high-speed, high-capacity disk storage subsystems; process them using custom high-speed, real-time, DSP microcomputer-based processors, hosted by UNIX workstations; communicate a variety of these signals between workstations over a high-speed fiber network; and output the processed signals at their intended destinations. Special-purpose systems - an HDTV workstation, and a custom DSP multiprocessor - will be integrated into the infrastructure to allow for real-time processing of high-resolution video sequences and images. The goal of this infrastructure is to build an environment in which the DSP algorithms, which will play a major role in the growth of multimedia technology, may be developed in such a fashion as to lead directly to real-time implementations of those algorithms. The companion part to this infrastructure is a program of research, based on this environment, being conducted in areas with immediate and urgent application to the problems facing multimedia information systems. This includes real-time video encoding, concentrating on the problems of inter- and intra-frame redundancy elimination. Advances in this area will lead to affordable applications in a host of areas such as video mail, video messaging, real-time video communications, and interactive learning and presentations.

Illinois Institute of Technology; Wai-Yip G Chan; CAREER: Audio-Visual Signal Compression Using Vector Quantization; (MIP-9502629); \$59,521; 36 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$112,041).

This research involves the development of efficient, reliable, and

real-time realizable coding algorithms and techniques for compressing and transporting audio-visual signals, targeting particularly the visual-telephone application. For a total channel bit-rate of 64 kb/s or less, our goal is to achieve superior grade-of-service by treating the compression of the audio (speech) and video signals as an integrated problem. Our thrust centers around the use of vector quantization techniques, and schemes that jointly optimize the signal analysis and quantization functions. We place particular emphasis on finding an efficient representation of the displacement ("motion") field for motion compensated prediction of the video signal. To facilitate personal wireless communications, adaptive coding schemes that exploit the prioritized transmission of encoded data are being explored. For compressing speech signals, variable rate coding schemes are being explored, and bits are being allocated between the speech and video signals to optimize the overall grade-of-service. Achieving this goal will provide the impetus to the development of the emergent Information Superhighway, forming part of an infrastructure supporting high performance computing.

Educational activities undertaken in conjunction with the above research include but are not exclusive to: the development of undergraduate digital signal processing laboratory courses; the introduction of graduate courses in the area of audio-visual information processing and communications; increasing the use of computer-based multimedia instruction in the classroom; and assisting the digitization of the library and networking of its databases.

University of Illinois; Yoram Bresler; *PYI: Statistical Techniques in Inverse Problems*; (MIP-9157377 A003); \$62,500; 12 months.

This research falls into four broad areas: image reconstruction, reconstruction of time-varying distributions, sensor array processing, and visualization of multiparameter data. In the area of image processing, the principal objective is to develop the theory and associated computational algorithms for superresolution image reconstruction from partial and noisy data using statistical models. For the second area, the goal is to develop optimum signal acquisition schemes subject to physical or economic constraints, and the associated efficient reconstruction algorithms, for imaging spatial data that is time varying during the acquisition process. In the area of sensor array processing, several issues are being addressed including the design of computationally efficient algorithms for the (sub)optimal solutions of model fitting problems, wideband source location, and imaging with sensor arrays. Finally, in the last area, the goal is to address the effective fusion, display, and visualization of multi-parameter spatially-related data, such as is acquired in multispectral, or multi-modality remote sensing and diagnostic imaging.

University of Illinois; Michael T Orchard; *NYI: Optimal Motion Compensation for Video Compression*; (MIP-9357823 A002); \$62,500; 12 months.

Efficient compression of video sequences should exploit the high interframe redundancy due to the smoothness of motion fields in typical scenes. Current video coding algorithms use very simplistic motion models, limiting the degree to which motion-induced redundancy can be exploited in video coding. This research is investigating improved methods for estimating motion in video sequences, and compensating for that motion to achieve increased coding efficiency. Methods are being considered which estimate motion at the encoder, requiring transmission of motion overhead, and those which estimate motion directly at the decoder. The goal is to provide a unified framework for these two approaches to motion estimation, and to develop hybrid algorithms taking advantage of the best features of both approaches.

University of Notre Dame; Ken D Sauer; *Model Based Tomography: A Comprehensive Approach to Iterative Image Reconstruction*; (MIP-9300560 A001); \$70,201; 12 months.

Low dosage transmission medical imaging, emission medical imaging, and nondestructive testing of materials are all examples of tomographic reconstruction problems which can benefit greatly from improved reconstruction techniques. This research introduces a strategy for the development of computationally efficient reconstruction algorithms which directly search for the statistically optimal fit to measured data. This approach exploits the wide availability of digital computation and storage to substantially improve reconstruction in demanding applications. The research has three essential components:

1. a general measurement system model which characterizes the physical measurement apparatus for both the transmission and emission problems, and can be extended to include nonlinear effects such as scattering;
2. a new class of computationally tractable image cross section models which preserve image detail while suppressing noise artifacts;
3. a fast numerical algorithm, known as Gauss-Seidel, which serves as a basis for efficient and accurate multiscale reconstruction methods.

Algorithmic techniques are being evaluated using data collected from the Nondestructive Evaluation Section at the Lawrence Livermore National Laboratories.

Johns Hopkins University; Jerry L Prince; *PFF: Three-Dimensional Image Processing*; (MIP-9350336 A001); \$100,000; 12 months.

This program has two major components: research and teaching in image processing, both having a focus on three-dimensional data. The research component is divided into two major thrusts.

1. The development of new methods for estimating motion in three dimensions from three-dimensional data sets. The first aim

involves optimizing data acquisition methods given a prior stochastic description of the motion and assuming the use of a certain optical flow method. The second aim explores the theory of three-dimensional vector tomography and its implementation using a magnetic resonance scanner.

2. The estimation of shape from three-dimensional data sets. This topic focuses on active surface methods: the definition of new active surfaces, the algorithms and their convergence properties, and the use of active surface methods for 3-D image registration.

The teaching component of this grant has two main objectives:

1. to develop a state-of-the-art signal and image processing laboratory, and
2. to develop new courses and textbooks on image reconstruction. The overall goal is to provide a comprehensive educational program at the cutting edge of signal and image processing

**Massachusetts Institute of Technology; Gregory W Wornell;
*CAREER: Theory and Application of Dispersive Multirate
Filterbanks and Wavelets*; (MIP-9502885); \$135,000; 36 months.**

While traditional signal representations developed out of multirate system and wavelet theory have been largely aimed at addressing signal analysis problems, the research component of this project is developing fundamentally new forms of these representations that are well suited to signal synthesis problems. In particular, non-localized, strongly-dispersive, and pseudorandom multirate signal expansions are being explored for a variety of applications ranging from multi-user wireless communications to broadband remote sensing. Wireless modems based on these techniques are being evaluated in a real-time DSP-based wireless communications laboratory testbed.

The educational component of the project encompasses curriculum development at both the undergraduate and graduate levels, and the mentoring of a diverse group of graduate students. This includes participation in the development of a new unified undergraduate introduction to the fields of communications, control and signal processing from a common foundation of basic signals, systems, and probability. Another focus is on re-engineering the graduate subject on detection, estimation, and stochastic processes, emphasizing new perspectives on random and deterministic signals, and on signal processing algorithms.

Michigan Technological University; Timothy J Schulz; CAREER: Image Recovery from Multiple Intensity Measurements with Applications to Electron Microscopy and Astronomy; (MIP-9501163, A001); \$137,795; 36 months.

The research component of this program has two major thrusts:

1. an investigation of the mathematical and computational aspects of the multiframe, bilinear image-recovery problem; and,
2. the development and application of solutions for problems in the areas of high-resolution transmission electron microscopy, and ground-and spaced-based astronomy.

Theoretical advances in this program are providing a unified approach to many imaging problems in which intensity data are recorded, including those for which satisfactory solutions have not been previously discovered. Collaborations with materials scientists on problems such as the study of atomic scale processes in solid batteries are providing exciting applications for which new methods developed in this program are being used to enhance the resolving and information-gathering powers of electron microscopes.

Through the teaching component, an educational program is being established with emphasis on:

1. the development of a multi-disciplinary laboratory where faculty and students from many departments jointly explore problems in remote sensing and imaging;
2. the development of multi-disciplinary courses in remote sensing imaging including students from many departments and tailored educational goals for each student -- ranging from theoretical foundations to specific applications; and
3. the use of a strong research-teaching exchange utilizing cutting-edge technology in the classroom and laboratory as a tool to excite, motivate, and retain enthusiasm in the students.

Stevens Institute of Technology; Sankar Basu; Multi-dimensional Nonseparable Subband Coding; (MIP-9322592 A001); \$44,132; 12 months.

This research is providing a parameterization of the entire family of multidimensional perfect reconstruction subband coding filter banks. The rationale for doing this lies in the practical application of efficient coding and compression of image (video) type signals. The parametric description of the family of multidimensional filter banks so obtained is being used to design perfect reconstruction filter banks having desirable frequency separation properties. The entire family of nonseparable multidimensional smooth wavelets resulting from iterations of these filter banks is also being completely parameterized in this way.

The techniques being adopted are from multi-dimensional system theory. The biorthogonal problems are approached by formulating them as matrix extension or matrix completion problems in the ring of multidimensional polynomials or stable proper rational functions; and the paraunitary problems as problems of describing, or equivalently, synthesizing multidimensional lossless systems. Similar strategies are used in the acausal formulation as well as in dealing with special classes of solutions such as the linear phase solutions.

Stevens Institute of Technology; Alan L Stewart, Roger S Pinkham; Self-Adjoint Operators and Models of Space-Variant Visual Acuity; (MIP-9405081 A001, A002); \$84,374; 12 months.

Current theories of visual processing are theories of local responses. Space-variant acuity concentrates on properties of the local receptive fields. Underlying this is the rationale that the visual system is approximately homogeneous within any small neighborhood. The response of the entire visual field is pieced together from models of local responses.

The theory of integral operators allows the experimental study of human acuity to be united with computational model in visual processing. Their use in this context leads to simpler and more intuitive proofs of key theorems which relate threshold assessment to eigenvalue problems. At the same time, everyday experimental concepts, such as threshold sensitivity, take on new elegance when placed within the theory of integral operators.

Brown University; Stuart A Geman, Donald E McClure, Basilis Gidas, Ulf Grenander; Mathematical and Computational Problems in Object Recognition; (DMS-9217655 A002); \$20,000; 12 months; (Joint support with the Robotics and Machine Intelligence Program, the Computational Mathematics Program, and the Statistics and Probability Program - Total Grant \$170,000).

The research program focuses on mathematical aspects of object recognition. There are two classes of problems. The first is the recognition of rigid objects positioned in a scene at arbitrary rotations, locations, and scales. A large repertoire of shapes is assumed known in advance, and the problem is to then devise computationally efficient algorithms for recognizing which, if any, of these objects are present in a given scene. Sequential and adaptive strategies will be explored, in which a sequence of image-based observations is made, with the choice of an observation depending upon the results of previous observations. There are close connections to coding theory, the sequential design of experiments, multi-armed bandit problems, the game of "twenty questions," and, of course, previous work in machine vision. The second class of problems is the recognition of nonrigid, or deformable, objects. Examples include handwritten numerals and various biological shapes, such as leaves, hands, organelles, etc. Here the issue of shape modeling appears to be central. An approach through deformable templates is proposed. Templates are prototypes which capture global characteristics, whereas deformations are random transformations, satisfying certain regularity constraints that

act upon templates to produce the possible presentations of an object. The proposed shape models suggest certain recognition algorithms which will be explored in a variety of application areas.

Texas Technological University; Frits H Ruymgaart; *Inverse Estimation Problems*; (DMS-9504485); \$40,000; 36 months; (Joint support with the Statistics Program - Total Grant \$105,000).

Inverse estimation is concerned with indirect curve estimation, where the curve of interest is to be recovered from observations that are subject to a random blur of a transformation of the curve. Inverse estimation arises when the object of interest can only be indirectly observed, and recovery of information about the object is required from the indirect measurements. Indirect measurement techniques, such as those used in medical imagery, are attractive because they are noninvasive. Perfect reconstruction of the image would, in principle, be possible if an unlimited number of uncorrupted measurements were available. However, one can only obtain finitely many data that are, moreover, corrupted by measurement errors.

The estimation problem can essentially be solved by inverting the transformation involved. Since this inverse is not in general continuous, the problem is typically ill-posed and regularization of the inverse is required. This research addresses two important questions: how restrictions due to limitations in time or space of the transformation relate to the unrestricted transformation; and, how recovery of irregular curves can be modified to avoid the Gibbs phenomenon.

George Mason University; David F Walnut; *Mathematical Sciences: New Results in Sampling and Wavelet Applications in Tomography*; (DMS-9500909); \$10,000; 36 months; (Joint support with the Applied Mathematics Program - Total Grant \$47,919).

This research has two parts. The first is applying new results in sampling theory to obtain high-resolution measurements of a signal from several low-resolution measurements of the same signal. A mathematical model of this problem (called multisensor deconvolution) is uniquely solvable but ill-posed, and difficult to solve numerically. The investigator has found Shannon-type sampling formulas on unions of regular lattices with incommensurate densities, which provide simple solutions to the multichannel deconvolution problem in special cases. These are being pursued. Several tomographic applications are being examined.

The second part of this program is developing "local" wavelet-based algorithms to recover edge features of an image from local Radon transform data. The following goals envisioned:

1. to find ways to recover locally density as well as edge features of an image, and to develop algorithms competitive with existing local tomography algorithms; and,
2. to identify the wavefront set of an image using wavelets, and to apply these techniques to the attenuated Radon transform.

University of Washington; Eve A Riskin; *NYI: Vector Quantization Codebook Processing and Organization*; (MIP-9257587 A003); \$62,500; 12 months.

New ways to use Vector Quantization (VQ) other than strictly for data compression are being investigated and applied to applications such as image processing, halftoning, progressive transmission, and immunity against communication channel noise. In many applications, both VQ and many image processing operations are applied to small subblocks of an image. The image processing step can be applied ahead of time to each vector in a VQ codebook, with the processed vectors stored along with the codebook. If the computational complexity of the VQ encoder is lower than that of the image processing step, this reduces the computational complexity. This approach is being applied to halftoning, edge detection, and histogram equalization. In a progressive transmission system, the received image is reconstructed as an increasingly better reproduction of the transmitted image as bits arrive. Ways to organize and order a VQ codebook so that it can be used for direct progressive transmission of full search VQ are being studied. In an ordered VQ, the VQ codeword index is correlated with the codeword location in the input space. This ordinal mapping feature of clustering codewords with similar indexes to obtain additional reproduction vectors for the decoder is being exploited. Extensions to progressive transmission of ordered VQ indexes over noisy communication channels are included.

University of Wisconsin; Truong Nguyen; *CAREER: Theory and Design of Filter Banks and Wavelets with Applications in Signal Conversion, Adaptation, Detection and Classification*; (MIP-9501589); \$120,219; 36 months.

High performance filter banks with high attenuation are needed in many applications such as audio compression algorithms, high-bandwidth high-resolution A/D converters, wideband signal detection systems and adaptive filtering algorithms.

The design of these filter banks are very difficult because of the nonlinear relations between the parameters and the objective function. This project is focused on the theory, structure, and design methods for Perfect-Reconstruction (PR)/Near-Perfect-Reconstruction (NPR) filter banks, and their applications in signal conversion, detection and adaptation. New structures and design methods are being investigated that will open up new classes of filter banks. They include biorthogonal cosine-modulated filter banks, infinite-impulse-response (IIR) cosine-modulated filter banks, and nonuniform filter banks. These filter banks are used in several applications such as high-performance A/D converters, nondestructive evaluation, echo cancellation, and adaptive noise cancellation systems. An undergraduate textbook is being written on the theory and design methods of filter bank wavelets. An internet site for the storage of design programs and coefficients of filter banks is being designed. A new course on time-frequency and time-scale analysis is under development.

Statistical Signal and Array Processing

Stanford University; Iain M Johnstone, David L Donoho; *Adaptive Estimation: New Tools, New Settings*; (DMS-9505151); \$70,000; 12 months; (Joint support with the Computational Mathematics Program and the Statistics Program - Total Grant \$190,000).

This research is developing statistical theory and computational tools in the general area of adaptive methods for representing and analyzing signals, images and other objects, and is showing how to tune them so they are noise-cognizant and stable. Underlying the approach are:

1. the idea of oracles, which know perfectly well how to ideally adapt representations;
2. the idea that the goal of adaptation in the presence of noisy data is to quantify how closely realizable procedures (which do not have privileged information about the object) can mimic an oracle; and,
3. the design of procedures coming as close as possible to the oracle.

The research is also developing methods for comparing different adaptation schemes by comparing oracles of different kinds, for example time-frequency oracles and time-scale oracles. This is an outgrowth of our earlier results on wavelets, where this approach was used to show that wavelets have a property of being nearly-ideally spatially adaptive. In addition a computational environment is being developed for implementing and systematically testing such approaches.

As a further outgrowth of the proposers' earlier work on wavelets, the project studies a number of improvements and extensions of wavelet shrinkage, for example in the directions of classification, confidence

bands, correlated data and selection of orthogonal bases. These efforts may have two spin-offs. First, some of the results may be stimulating and/or useful to the community of "inventors of adaptive procedures" in signal, image, speech, and time/frequency, and related communities. Second, the theoretical work may stimulate statisticians to take more interest in making further contributions in such directions.

University of California-Davis; William A Gardner; *Programmable Blind Adaptive Multivariate Filtering*; (MIP-9412732); \$185,650; 36 months.

In wireless communications, including cellular communication systems, spread spectrum overlay systems, and signals intelligence applications, the degradation caused by rapidly time-varying multipath and unknown cochannel interference can be reduced by adaptive spatial filtering using antenna arrays. This research is pursuing a flexible framework for adapting a spatial filter without using a training signal, array calibration data, or knowledge of spatial characteristics of the desired or interfering signals. It is derived from two conceptually different perspectives: canonical correlation analysis from multivariate statistics, and the conditional maximum likelihood problem for a data-derived training signal obtained from a user-programmable transformation of the received data. The transformation can be programmed by the user to sort and separate the received signals on the basis of the differing degrees to which they exhibit one or more user-selected statistical properties. The performance of the new method is being investigated analytically and by computer simulations to quantify its capabilities of signal separation, multipath mitigation, and interference rejection. Preliminary results suggest that the new method can converge very quickly to yield signal estimates that are comparable to those obtained by the Mean Squared Error (MSE) method that uses known training

signals. More generally, the objective of this research is to develop and evaluate a new and general approach to programmable blind adaptive multivariate filtering that has broad application to sensor array processing in addition to wireless communications.

University of Florida; Jian Li; *NYI: SAR Image Formation and Processing Techniques for Environmental Monitoring*; (MIP-9457388 A001); \$26,000; 12 months.

This research shows the advantages of using Synthetic Aperture Radar (SAR) to detect, analyze, and quantify environmental changes. There are three foci.

1. Improve SAR image formation with a phased array airborne or spaceborne radar. Efficient SAR image formation techniques are being developed by appropriately designing the transmitted waveforms of the phased array radar. The trade-offs between nonparametric and parametric techniques are being studied.
2. Research on the SAR image understanding and ground truth evaluation. Statistical clustering algorithms and various feature extraction schemes that adequately incorporate electromagnetic phenomena are being developed and evaluated.
3. Research on the detection, analysis, and quantification of environmental changes through repeated SAR imaging of critical regions including the environmentally fragile Florida wetlands. Change detection techniques are being developed to quantify and document subtle environmental changes from these images.

Purdue University; Michael D Zoltowski; *Closed-Form Angle Estimation with Circular Arrays/Apertures for Mobile/Cellular Communications and Surveillance Radar*; (MIP-9320890 A001); \$24,223.

The digital communications industry is currently investing enormous resources towards the development and experimental verification of prototype antenna arrays to be deployed on mobile communication vehicles, including the commercial automobile of the future, as a means of discriminating amongst signals co-located in frequency based on their respective spatial locations. Given the small aperture on a mobile communications unit, the Uniform Circular Array (UCA) geometry is ideal due to its rotational invariance with respect to azimuth. This research is based on a recent development of a simple, closed form algorithm, UCA-ESPRIT, for use in conjunction with a UCA that provides automatically paired source azimuth and elevation angle estimates. To date, the algorithms for 2D arrival-angle estimation have required expensive spectral searches, iterative solutions to multidimensional optimization problems, or ad-hoc schemes for pairing direction cosine estimates with

respect to each of a number of different array axes. UCA-ESPRIT is fundamentally different from ESPRIT in that it is not based on a displacement invariance array structure but rather is based on phase mode excitation and hinges on a recursive relationship between Bessel functions.

A theoretical performance analysis of UCA-ESPRIT is being conducted. This proves extremely useful for predicting its performance in a mobile communications environment. Novel strategies for incorporating mutual coupling effects are also being developed. The real world performance of UCA-ESPRIT is being assessed with experimental data from a prototype circular antenna array currently being built at the Polytechnic University of Madrid for mobile sea communications with the INMARSAT satellite system. Adaptations of UCA-ESPRIT for filled circular arrays are being developed.

State University of New York - Stony Brook; Petar M Djuric; *Bayesian Solutions to Model Selection, Parameter Estimation, and Spectral Analysis*; (MIP-9506743); \$119,098; 24 months.

Model selection, parameter estimation, and spectral analysis are three important areas in statistical signal processing. This research explores some difficult and unresolved problems in these disciplines by exploiting Bayesian theory. Topics of interest include the derivation of model selection rules based on asymptotic assumptions and their applications to problems in array processing, rank determination in time series analysis, and segmentation of vector fields; analysis of transient signals and parameter estimation of highly nonlinear models such as threshold signal models and bilinear models; and, Bayesian spectral analysis of nonstationary signals.

This effort primarily consists of three equally important components:

1. a theoretical investigation into these problems that leads to an improved understanding of various signal models and concepts;
2. the practical application of the solutions, which includes automatic segmentation of medical images and the processing of single channel patch clamp currents; and,
3. the student's use of a practical exposition into the versatility of Bayesian inference and its applicability for solving a wide range of signal processing problems.

University of North Carolina; Jianqing Fan, James S Marron; *Mathematical Sciences: Processing Massive Noisy Data with Hidden Structure*; (DMS-9504414); \$60,000; 36 months; (Joint support with the Statistics Program - Total Grant \$135,000).

Many scientific disciplines depend in some way on extracting structural information from noisy data. Fields ranging from processing noisy images and evaluating business marketing, to analyzing survival data and forecasting economic climates, which are universes apart in their backgrounds, nevertheless share the common problem of drawing conclusions via the processing of noisy signals. Such problems may be abstracted as statistical function estimation problems, and can be analyzed by various techniques in this project.

The objective of this research is to develop and evaluate flexible statistical modeling techniques. Specifically, to develop new statistical methodologies and to investigate their foundational properties for flexible statistical modeling in processing high-dimensional data, nonparametric confidence intervals, mode detection, and signal processing. These techniques can be applied in the Federal Strategic Areas such as monitoring environmental and global changes via processing massive collected data, where informative structures can hardly be detected by traditional approaches, and building statistical modeling for economic and business activities in the civil infrastructure. The investigators will take advantage of modern computing facilities and use statistical knowledge to avoid unnecessary data mining, hence, significantly reducing the data processing time.

field. Examples of imaging array systems occur in radio astronomy, sonar, and microwave and ultrasound imaging.

This research is developing new results for array design and associated signal processing, based on recent developments on the characterization of array performance in linear imaging. Active imaging systems (e.g., ultrasound imaging arrays) are a particular focus of this investigation, although some of the work also addresses passive arrays. Using the idea of the "coarray", this work considers how to deploy array elements (and associated hardware) in the most efficient way to obtain large array apertures and high resolutions. The results allow minimum redundancy active arrays and minimum complexity active arrays to be specified. This research on arrays includes a study of their characteristics under real operating conditions and extends to some experimental work with an acoustic array system.

Lehigh University; Rick S Blum; *RIA: Distributed Signal Detection in Uncertain Environments*; (MIP-9211298 A002); \$10,000.

Distributing multiple sensors over some region for the purpose of detecting a signal in noise is becoming increasingly attractive. The majority of research work concerning the design of such schemes has focused on signal detection problems where a complete observation model which categorizes the environment under consideration is known. This research deals with distributed detection schemes for cases where a complete observation model is not available due to incomplete knowledge of the operating environment. This is an extremely practical case which has applications in air traffic control, radar weather monitoring, and other radar and sonar systems. Robustness results with alternative hypothesis which are not simple, finite observation sample sizes, dependent observations, and non-additive observations are being sought. The plan to construct these schemes focuses on applying the theory of minimax robust statistics. Cases with nonstationary, possibly non-Gaussian background environments and possibly dependent observations are also being investigated.

University of Pennsylvania; Saleem A Kassam; *Arrays for High Resolution Imaging and Efficient Digital Filtering*; (MIP-9321856 A001, A002); \$69,462; 12 months.

Many imaging systems use arrays of individual elements to sense the propagating field produced by an object. Through signal processing techniques such as beamforming, an image of the object may then be formed. In addition to such passive arrays, imaging systems may use active arrays to both illuminate an object with a radiated field and to record the reflected

University of Texas at Austin; Guanghan Xu, Dim-Lee Kwong; *Development of Advanced Signal Processing Algorithms for On-Line Temperature Profile Measurement in Semiconductor Manufacturing*; (MIP-9400732 A001, A002); \$989,203; 12 months.

This research is directed towards the development, implementation, and demonstration of advanced model-based signal processing algorithms for real-time measurement of wafer temperature profile in Rapid Thermal Processing (RTP). RTP cluster tools are strategically important for submicron semiconductor manufacturing because of trends towards reduced thermal budget and tightened process control requirements on large diameter silicon wafers. Despite its significant advantages, commercial versions of RTP modules for various chemical vapor deposition applications have not been available.

This project is developing advanced model-based signal processing algorithms, which when coupled with the acoustic thermometry and acoustic/pyrometer approaches, accurately measure the wafer temperature profile at fast acquisition rates and with a minimum number of sensors. This entails algorithm development, implementation, and validation using real data from commercial RTP tools that are available at the University of Texas at Austin and SEMATECH.

University of Texas at Austin; Guanghan Xu; *CAREER: Development and Implementation of Antenna Array Processing Techniques for Wireless Communications*; (MIP-9502695); \$88,857; 36 months.

Array signal processing techniques were traditionally limited to military applications. Researchers recently found that these techniques could be applied to significantly expand channel capacity and improve quality of wireless communication systems through exploitation of spatial diversity. Despite significant research activities in array processing during the last decade, conversion of military technology to commercial technology has not been cost-effective. A significant amount of research and development effort is required to realize these techniques in a wireless communication system. This research is developing innovative antenna array processing techniques and fast implementation schemes for various wireless communication systems. The program scope is not limited to algorithm development, theoretical analysis, and simulation studies; our ultimate objective is to implement the resulting algorithms in real hardware.

Industries specializing in digital signal processing and its applications in telecommunications have experienced rapid growth due to a high demand for information access and processing. To respond to such a growth, new courses in these areas and innovative

teaching methodologies must be developed so students can quickly adapt to various challenges. Plans include:

1. the development of several new courses in these areas including undergraduate laboratory courses;
2. the continued involvement of undergraduate students in industry sponsored projects;
3. the design of more computer projects;
4. the creation of opportunities for minority students and students with disabilities; and,
5. the development of a multimedia teaching method drawing on advanced technology, e.g., live audio, video, and graphical illustrations and real hardware demonstrations.

Brigham Young University; A. L. Swindlehurst; *Analysis and Development of Algorithms for Antenna Array Based Communications Systems*; (MIP-9408154); \$50,388; 12 months.

Communication systems employing antenna arrays are required in situations where multiple co-channel signals are present simultaneously. The spatial discrimination provided by multiple antenna elements allows spectrally overlapping signals to be individually extracted. Such systems have typically been proposed in the context of military applications, but important commercial applications have recently gained attention. It has been proposed that antenna arrays be used in land-based mobile radio systems to provide enhanced spatial discrimination and hence increased capacity. Such arrays could potentially provide frequency reuse in adjacent cells, or within a communication cell itself.

This research is examining signal copy, interference cancellation, and source localization using antenna arrays, with emphasis on scenarios that might be encountered in mobile radio communications (e.g., multipath fading environments, signal formats such as analog FM IS-3 AMPS, GSM, IS-54 offset QPSK, IS-95 CDMA, etc.). This research focuses on the development and analysis of algorithms that exploit all available temporal and spatial information. Conventional techniques for Direction Finding (DF) and Signal Copy (SC) rely on either spatial or temporal structure in the data, but typically not both. Three different techniques that exploit both spatial and temporal signal structure for improved DF/SC are being studied:

1. iterative blind-least-squares;
2. decision directed algorithms array-based equalization of multipath channels; and
3. optimal joint DF/SC with parametric array uncertainty.

Finally, special emphasis is being placed on the use of real mobile cellular radio data to test the effectiveness of the algorithms developed.

University of Virginia; Georgios B Giannakis, Michail K Tsatsanis; *Estimation and Equalization of Time-Varying Channels*; (MIP-9424305); \$114,428; 24 months.

InterSymbol Interference (ISI) is caused by multipath effects or bandwidth constraints and presents a major obstacle in modern high-speed digital transmission. When the communication link is fading with time, it is essential for the equalizer to be able to follow the

channel's time variations. While several adaptive and blind methods have been proposed for time-invariant or slowly fading channels, relatively little attention has been given to the rapidly fading case.

The goal of this research is to use basis expansion ideas as a tool for the equalization of rapidly varying channels. In this framework, each time-varying coefficient is expanded into a set of basis sequences, and the expansion parameters are estimated using novel adaptive and decision-directed schemes. Next, basis expansions are combined with time-varying correlations and higher-order statistics to address time-varying, blind equalization in a rapidly fading environment. Finally, based on physical intuition and experimental data, the selection of the "best" basis which captures optimally the channel's dynamics is being addressed. Complex exponentials are being investigated for the mobile radio channel; wavelet expansions are being considered for channels with abrupt changes, and are being compared with stochastic channel modeling approaches.

Other

Stanford University; Arogyaswami J Paulraj; *International Conference on Communication, Computing, Control and Signal Processing, Stanford, California, June 22-26, 1995; (ECS-9526114); \$5,000; 2 months; (Joint support with the Systems Theory Program - Total Grant \$10,085).*

Traditional barriers across disciplines such as Communications, Computing, Control, and Signal Processing are fast disappearing. New technological challenges require concepts and techniques from several disciplines. Interdisciplinary approaches are being increasingly developed to tackle challenging problems in seemingly disparate areas. Driven by fundamentally new applications and approaches, the coming century needs even further convergence and cross-fertilization of ideas. The objective of the conference is to highlight some recent advances made toward developing an integrated approach to problem solving at various levels, and to communicate some of the exciting questions and challenges that remain unsolved. The conference stimulated discussions at various levels:

1. mathematical foundations;
2. algorithms (statistical and computational efficiency, numerical issues, and new paradigms); and,
3. architectures (parallel systems, hardwired, reconfigurable and programmable).

University of California-Irvine; Kai-Yeung Siu; *NYI: Analysis and Design of Artificial Neural Networks; (MIP-9357553 A002); \$31,250; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$62,500).*

Artificial neural networks present a new model for massively parallel computation and a paradigm for solving large scale optimization problems. This research is exploring the advantages of neural network-based models over conventional models for computation, and a novel design of neuromorphic computing architectures for applications in signal and image processing. A theoretical framework is being established to derive tight tradeoffs between the number of elements and the number of layers in neural networks. The results should answer some of the key open questions in the analysis of neural networks using classical mathematical tools such as rational approximation techniques and harmonic analysis.

University of Maryland; K. J. Ray Liu; *NYI: High Performance Computing for Signal Processing; (MIP-9457397 A001); \$31,250; 12 months; (Joint support with the Systems Prototyping and Fabrications Program - Total Grant \$62,500).*

There are three major architectural models used in high-performance signal/image processing:

1. VLSI- signal processing - high- throughput VLSI architectures for low-cost application-specific implementations used in applications such as communication systems, speech, video/HDTV, and radar;
2. parallel signal processing on massively parallel computers - parallel algorithms for complex signal/imaging systems used in computer vision, medical imaging, and the processing of vast amounts of data in deep space exploration;
3. distributed signal processing on high-speed networks - used in applications such as document image processing, multimedia, automatic signal processing in manufacturing, and medical signal/image processing.

This research focuses on the development of efficient algorithms and architectures for each architectural model. Comparative studies of the advantages and disadvantages of these different computing

schemes are being conducted. The goal is to investigate which signal/image processing problems can be carried out optimally under different computing and communication schemes.

University of Minnesota; Keshab K Parhi; *NYI: Dedicated VLSI Digital Signal and Image Processors*; (MIP-9258670 A003); \$62,500; 12 months.

Research efforts are being directed towards the design of dedicated, high-performance digital signal and image processors. The emphasis is on real-time processing, where samples are processed as they are received from the source, as opposed to being stored in buffers and then processed in batch. Design of algorithm topologies for recursive signal processing algorithms were once considered a major challenge. Using the relaxed look-ahead technique, new concurrent algorithms and topologies for adaptive LMS and lattice filters, cascade and lattice recursive digital filters, and predictive speech and image coders have been developed. Design of concurrent topologies for wave digital filters, decision-feedback equalizers, and adaptive differential vector quantizers are being pursued. The decoding speed in Huffman and arithmetic coders (used for lossless compression) is limited due to the feedback. For the Huffman decoder, the codeword length multiplicity constraint is being exploited to design codes where multiple bits can be simultaneously decoded in parallel. The performance of these decoders is further improved by the use of conditional coding. Novel approaches for design of parallel arithmetic coders are also being pursued.

State University of New York - Stony Brook; Michael M Green; *NYI: Improved Circuit Simulation Using Results from Circuit Theory*; (MIP-9457387 A001); \$31,250; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$62,500).

This research is applying results in the area of nonlinear circuit theory to enhance the simulation of analog circuits. Improvements to the continuation methods of solving dc operating points of circuits are being made to guarantee that all circuit operating points are found during a single analysis, and are being applied to sensitivity analysis of circuits. Erroneous models are thought to be a major source of convergence problems and erroneous results in circuit simulation. Another enhancement to circuit simulation includes checking the accuracy of transistor models by verifying that all models satisfy passivity and the no-gain condition.

Brown University; Harvey F Silverman; *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A002, A003, A004); \$60,000; 24 months; (Joint support with the Systems Prototyping and Fabrications Program, Experimental Systems Program, Microelectronic Systems Architecture Program, Design, Tools and Test Program - Total Grant \$674,601).

This is a collaboration between Brown and Rutgers Universities to build a large microphone array with associated processors for beam forming. Applications are direction finding, echo cancellation, and speaker differentiation in telconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

Brown University; Harvey F Silverman; *Parallel Architectures for Speech Recognition: Nonlinear Optimization of Expectation-Maximization (EM) Training of Hidden Markov Models (HMMs) in a Reconfigurable Environment*; (MIP-9509505); \$50,815; 12 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$187,251).

This research focuses on one of the strategic areas of national concern, that of high performance computing. It involves the development, construction and testing of new architectures for high-speed computing. The idea is to combine general-purpose, RISC-based processing nodes, each with multiple field-programmable-logic-array(FPLA) based coprocessor systems. This kind of system has the hardware properties of a general-purpose computer, but the advantage of performance more akin to that of a special-purpose engine. In particular, the 20-node Armstrong III system has been built, is operational, and provides nearly two-orders of magnitude improvement over current advanced workstations. The hardware

system, combined with its configuration compiler, a program that automatically translates a C function subroutine to both machine code and to the hardware design of the reconfigurable coprocessors, are the basic building blocks for this class of architectures.

This project is unique in that not only has the hardware/software system been built, but it is also being tested on real applications. The training of a modern hidden markov model based speech recognition system requires hundreds of hours, even with recent variants developed at Brown and elsewhere. The reduction of this training time to ten or so minutes allows progress in this area to be made at a faster rate and/or that expensive nonlinear optimization techniques may now be applied to the problem. Also, this means that data from a less cumbersome sensor system (the microphone array systems in place and being developed at Brown) can be suitably incorporated into a more robust speech recognition system.

William Marsh Rice University; Don H Johnson; *Databases for Signal Processing Research*; (MIP-9301646 A003, A004); \$43,302; 12 months.

Rice University is establishing a two-part database composed of sampled signals and signal processing software with the Signal Processing Society serving as the gatekeeper. This database is accessible through the InterNet at no charge to users. Data and programs are being solicited from university, industrial, and military sources. The data provides a needed testbed for evaluation of signal processing algorithms; the software provides the signal processing community with state-of-the-art algorithms and simulation systems.

Washington State University; Roberto H Bamberger; *A Prototype Distance Learning Computer Laboratory for Exploratory, Asynchronous Learning*; (REC-9552985); \$15,000; 12 months; (Joint support with the Course and Curriculum Program, and the Applications of Advanced Technologies Program - Total Grant \$49,676).

This project is developing and field testing three key components of a distance learning computer laboratory:

1. an interpersonal communication tool which combines elements from email, video-conferencing, listserv/listproc, usenet news, WAIS, gopher, and WWW;
2. network based, interactive, media-rich, digital collections or curricular material; and,
3. authoring agents for the seamless integration of the interpersonal communication systems and interactive digital collections.

Deliverables include the prototype systems, raw data for assessment of the impact of the tools on educational delivery, and a preliminary assessment study. To fully demonstrate the need and benefits of distance learning computer laboratories, the prototype system is being developed in the context of a course on digital image processing.

Experimental Systems

Dr. Michael J. Foster, Program Director
(703) 306-1936 mfoster@nsf.gov

The Program

The Experimental Systems program supports research projects that involve building, evaluating, and experimenting with a computer or information-processing system. These are goal-oriented projects generally undertaken by teams of designers, builders, and users. The building of the system must itself represent a major intellectual effort, and offer advances in our understanding of information systems architecture. A system supported by the Experimental Systems program will usually include both hardware and software components.

Research on information processing systems involves interaction among diverse elements such as hardware architectures, computational models, compilers, operating systems, applications, performance evaluation tools, and user interfaces. Building and evaluating real experimental systems is the only way to understand these interactions in large systems; other techniques, such as simulation and analysis, have only limited uses in understanding the system issues in such a complex environment. Software simulators, for instance, do not provide the computing speed needed for large experiments, nor the needed performance incentives for porting large application systems for experimentation. Without real experimental systems, important areas of information systems architectures cannot advance.

A successful proposal to the Experimental Systems program should demonstrate the feasibility and utility of the project. Feasibility can be shown by describing prior proof-of-concept prototypes or simulation studies that indicate that the proposed system can be built and will meet its design goals. Utility can be shown by demonstrating that building the system will provide substantial advances in computer system architecture, or that the system is inherently useful. Details of the measurement and evaluation procedures that will demonstrate the benefits of the system in an application should be given in the proposal.

The system to be built must be novel in some way, and the impact of the novel aspects of the system upon its architecture must be evaluated during the course of the research. To justify construction the new system must be potentially superior to existing systems in the chosen application area. Ideally, building the system would provide new knowledge of systems architecture, open up new application areas, and/or contribute to our knowledge about system building techniques. An appropriate project might be a system built using a new architecture or technology, which addresses an application in a new way. An inappropriate project would be one in which the research uses, simply as a platform, a special purpose machine whose design, fabrication, and evaluation are straightforward. The novel aspects of an experimental system may fall into several different areas; the system might feature application of a new technology, new architecture, or new techniques for performance measurement and evaluation to a computationally stressing problem. Examples of technological innovation are massively parallel analog systems, or applications of opto-electronics. Architectural innovations might include new parallel I/O structures, hardware-software codesign, or limited modifications to commodity processors. New evaluation techniques might include instrumentation for performance evaluation or debugging. These innovations might be applied to produce high-performance computers, intelligent sensors, or signal processing architectures, for example. A list of projects currently supported under the Experimental Systems program can be found in the Microelectronic Information Processing Systems Division Summary of Awards.

To justify support under this program, a proposal should show that system building is necessary for answering significant and timely research questions. The research issues should be such that the best way to address them is to build the proposed system and measure its performance. Building for its own sake is discouraged; analysis and simulation should be performed in sufficient detail before a proposal is sent to the Experimental Systems program. Furthermore, off-the-shelf hardware should be employed in the building stage whenever the research goals do not require custom construction.

By encouraging the design, construction, test, and evaluation of novel information processing systems, NSF hopes to achieve several goals:

- * Settle major research issues and add to fundamental knowledge in information processing;
- * Guide university research in computer science and engineering toward meaningful problems of industrial interest;

- * Strengthen the system-building expertise in our research institutions;
- * Educate a new generation of researchers in experimental systems research.

Potential applicants are encouraged to discuss their research ideas with the program director prior to formal submission.

Initiatives and Opportunities

During the 1995 Fiscal Year, the Experimental Systems program participated in several initiatives. These initiatives were intended to encourage research directions that had been identified as important by the research community and that crossed program lines.

HIGH PERFORMANCE STORAGE SYSTEMS AND WIDE-BAND I/O

This initiative resulted from recommendations developed by the NSF Workshop on High Performance Memory Systems, and by the Information Infrastructure Technology and Applications Task Force. Proposals that addressed problems of accessing large data sets at a high rate over networks were especially encouraged. Topics of interest included the following:

- * The development of new memory systems that take advantage of emerging storage technologies;
- * New techniques for organizing cache memory and other buffering schemes to alleviate memory and network latency;
- * Partitioning of systems to reduce data movement;
- * Reliability and fault-tolerance of new memory systems.

RAPID PROTOTYPING: VIRTUAL AND PHYSICAL

Emerging national needs and recent technological advances point to rapid prototyping as an essential area of research in the strategic area of advanced manufacturing. The goal of research in rapid prototyping is to develop and integrate the tools and technologies needed for rapid and efficient design and manufacturing of products, processes, and systems. The result will be reduced product delivery times to meet dynamic market requirements.

RESEARCH IN NEW GENERATION OPERATING SYSTEMS ARCHITECTURES

To meet the diverse needs of applications on diverse computing platforms, today's operating systems have been growing in size and complexity, incurring more overhead in providing required services. There is a need for a systematic evaluation of the abstractions needed for the design and implementation of operating systems for the current and emerging generations of system architectures and applications.

Several research issues related to this initiative are of interest to the Experimental Systems program. These generally involve architectural support for operating systems on networks of workstations or on collections of supercomputers connected by gigabit networks.

Awards

Storage Hierarchies and Input/Output Systems

University of Southern California; Michel Dubois, Rafael H Saavedra, Massoud Pedram, Peter B Danzig; *The USC Multiprocessor Testbed: A Testbed for Scalable Shared-Memory Systems*; (MIP-9223812 A001); \$558,021; 14 months.

A testbed for experimenting with memory hierarchies in multiprocessors is being supported. A processor node in the testbed contains cache and memory system controllers made from field-programmable gate arrays. To experiment with a memory control mechanism or coherency technique, the investigators program the gate arrays to implement the mechanism. For software support of experimental techniques, the GNU-C compiler is being modified to generate appropriate code, such as non-blocking prefetches, and the Mach microkernel is being ported to provide thread scheduling.

Massachusetts Institute of Technology; John L Wyatt, Ichiro Masaki; *Cost-Effective Hybrid Vision System for Intelligent Highway Applications*; (MIP-9423221); \$400,000; 12 months.

The goal of this project is a new cost-effective architecture for machine vision, which will be evaluated for intelligent highway applications. Components of three smart image sensors are being developed and integrated into a heterogeneous vision system. The smart sensors use analog, digital, and mixed signal techniques to perform 3 dimensional measurement for adaptive cruise control, lane detection, and time-to-collision measurements. All of these tasks are intended for machine vision systems within intelligent vehicles, and will be tested in intelligent vehicle applications within the later stages of the award. MIT is performing basic research on the first versions of all of these cameras, while later development, integration, and testing will be carried out by industrial partners.

State University of New York - Stony Brook; Tzi-Cker Chiueh; *CAREER: Parallel I/O for 3D Volume Visualization*; (MIP-9502067); \$165,000; 36 months.

This project is applying parallel I/O technology to reduce I/O delay in volume visualization. Volume images that arise from such applications as computer tomography and 3D ultrasound are too large to fit within random access frame buffers. As a result, disk I/O is in the critical path for volume rendering. This project is exploring the use of a new parallel volume

rendering technique that traverses 3D data sets in an order that eliminates I/O conflicts in a parallel disk array. Consequently, the entire bandwidth of the disk array is available to the rendering application. In this project, a high performance volume visualization server using this idea is being designed, implemented with off-the-shelf components, and evaluated.

A supplement supporting Research Experiences for Undergraduates will permit advanced undergraduate students to contribute to the project by implementing and measuring I/O intensive applications.

Carnegie-Mellon University; Takeo Kanade, Andrew Gruss, Dean Pomerleau, L. Richard Carley; *ALVINN-On-A-Chip: A Computational Sensor for Road Following*; (MIP-9305494 A001, A002); \$316,740; 12 months.

This project is building and deploying an intelligent imaging sensor for road following. The sensor generates the heading information required to steer a robotic vehicle by watching the road. On-chip processing is performed by a neural network trained to drive autonomously on public highways. The circuitry which performs the neural computations is integrated with a photosensor array in order to directly sense road-image information. The photosensor array includes analog signal processing in each cell and binary optics for better photon statistics, decreased transducer size, and less interference.

Brown University; Harvey F Silverman; *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A002, A003, A004); \$314,601; 24 months; (Joint support with the Systems Prototyping and Fabrications Program, the Circuits and Signal Processing Program, the Microelectronic Systems Architecture Program, and the Design, Tools and Test Program - Total Grant \$674,601).

This is a joint project between Brown and Rutgers to build a large microphone array with associated processors for beam forming. Applications are direction finding, echo cancellation, and speaker differentiation in teleconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

University of Utah; Lee A Hollaar, Ellen M Riloff;

Implementation and Evaluation of a Parallel Text Searcher for Very Large Databases; (MIP-9023174 A007); \$18,849.

This project concerns the application of the Utah Retrieval System Architecture to very large databases of full-text documents. This effort involves the development of a medium-scale (4 to 10 GBytes) parallel backend search server using augmented RISC processors as the searching engines. Data is being gathered and analyzed to determine if the existence of a high-speed search server changes the complexity and arrival rate of queries by real users (law students). In addition, a suitable partitioning of functionality such that remote users can be supported by such a searching engine over medium-speed networks (such as ISDN) is

being studied. The researchers are also examining how the system can be reconfigured to deal with disk and searcher failures.

University of Virginia; William A Wulf, Jack W Davidson, James H Aylor; Implementation of High Bandwidth Memory Systems; (MIP-9307626 A001); \$400,000; 12 months.

This project is building and measuring experimental memory systems that match the high data rates of processors with the low random access data rates of memory parts. The goal is to detect streams of memory references at compile time and use a smart memory controller to prefetch the streams at run time. The memory controller can use features of the memory system such as page mode, nibble mode, or Rambus to maximize the data rates of the memory parts. It then buffers the streams until the processor asks for the data. The project includes compiler research as well as research into the architecture of the memory controllers. This is similar to the vectorization efforts on such machines as the Cray supercomputers, but more general since a smart memory controller can be designed for any combination of processor speeds, memory features, and program characteristics.

General Purpose Computing

Stanford University; Michael J Flynn, Bruce A Wooley, S. Simon Wong, Giovanni De Micheli, Fabian W Pease; Sub-Nanosecond Arithmetic II; (MIP-9313701 A001); \$320,000; 12 months.

This project is attempting to speed up computer arithmetic by several orders of magnitude using a combination of algorithmic, circuit, and packaging techniques. CAD tools to automate the application of these techniques are also being developed under the project. Specific research problems include the development of a package capable of passing large numbers of signals with 100 picosecond rise times, and the integration of wave pipelined data paths into an overall system.

to: study the optimal division of responsibilities among the hardware, compiler, and operating system to maintain cache coherence in scalable share-memory multiprocessors; design algorithms and hardware to effectively support multiprogramming of parallel programs in scalable shared-memory multiprocessors; and optimize the management of memory hierarchies and the interaction of the operating system with the architecture.

University of Illinois; Josep Torrellas; NYI: Increasing the Performance of Scalable Shared-Memory Multiprocessors; (MIP-9457436 A001); \$75,000; 12 months.

Scalable shared-memory machines are a promising way of attaining large-scale multiprocessing without surrendering much programmability. Achieving high performance from these machines, however, is challenging because many complex architecture and architecture-software implementation issues that have been only partially studied considerably impact the performance of the machines. The objective of this research is to contribute in three areas to help make shared-memory multiprocessors the preferred source of computing power. The three thrusts of this project are

Massachusetts Institute of Technology; Anant Agarwal; PYI: Automatic Locality Management in Scalable Multiprocessors; (MIP-9157393 A004); \$62,500, 12 months.

Parallel computers can be made both scalable and easily programmable through architectures that exploit and automatically manage communication locality. The goal of this research is to discover and to evaluate techniques for automatic locality management in scalable multiprocessors. As the vehicle for this research, an experimental parallel machine called the Alewife is being implemented. Alewife employs techniques for:

1. communication latency minimization, using scalable coherent caches and software partitioning and placement of programs, and
2. communication latency tolerance, using a new rapid-context-switching processor architecture.

Alewife implements a new protocol called "limitless directories" for scalable cache coherence. This scheme uses a combination of hardware and software techniques to realize the performance of a full-map directory with the memory overhead of a limited directory. A rapid-context-switching processor called Sparcle is also being designed. Sparcle can switch in about 10 cycles to another thread

when it suffers a cache miss that requires service over the interconnection network.

The major goal for this grant period is to get a small prototype Alewife system operational, including the hardware as well as the entire software system.

Massachusetts Institute of Technology; Anant Agarwal; *Protection and Translation in Multimodel Multiprocessors: The MIP FUGU Workstation*; (MIP-9504399); \$563,735; 24 months; (Joint support with the Microelectronic Systems Architecture Program - Total Grant \$800,000).

This project is extending the Alewife multiprocessor to explore two aspects of future scalable multiprocessors. One aspect is protection. Protection is required on accesses to hardware resources, with as little software overhead as possible. For example, when a message sender and receiver are part of the same application, they should be able to directly access network hardware. The other aspect is address translation, so that each application can have its own address space. All processors executing an application must have the same address space, however. To ensure this, memory mapping hardware is placed at each node of the multiprocessor, and must be kept coherent.

The FUGU system provides three modes of interprocessor communication. Shared memory is provided by hardware synthesized messages among nodes, which maintain cache coherence. Short messages can be sent directly by performing loads and stores to I/O registers that are in each node's address space. Longer messages are managed by a DMA engine that is separate from the processor.

Translation and protection for shared memory and for short messages are provided by hardware enhancements that maintain coherence in the translation lookaside buffers and ensure that messages are received only by the processes that they are addressed to. For long messages, only minimal support for address translation is provided by the hardware since the need for such translation is likely to be rare. Any needed translation will be provided by software.

Princeton University; Kai Li, Margaret R Martonosi, Douglas W Clark, Edward W Felten, Richard J Lipton; *SHRIMP: Architectural and Systems Support for Inexpensive, High-Performance Multicomputers*; (MIP-9420653); \$399,911; 12 months.

This project is building a high-performance multiprocessor from commodity desktop computer systems and off-the-shelf interconnects. Commercial Intel Pentium workstation boards, each with attached memory, disk, and I/O, are attached to a Paragon backplane. Communication uses a new mechanism called virtual memory-mapped communication, which

disguises interprocessor communication as write operations to memory. The node interface maps physical pages in the memories of individual nodes to each other, so that a write to one mapped page results in messages to other nodes that share the mapped page. The operating systems on the individual nodes use their ordinary virtual memory mechanism to support virtual page mapping. In addition to this word-by-word communication, DMA transfers are available, with control registers located in the address space of individual processes. This allows high bandwidth communication that maintains user-level protection. Research to be addressed in the project includes the achievement of high-bandwidth low-latency communication between processes, the structure of an I/O system supported by the new communication mechanism, and performance evaluation of the resulting system.

New York University; Allan Gottlieb; *Evaluating the NYU Ultracomputer*; (MIP-9303014 A002); \$351,079; 12 months.

This is a project to characterize and model the performance of a scalable shared memory computer. Ultracomputer uses a multistage interconnection network with hardware combining to provide high-bandwidth scalable connections between processors and memory. Ultra III, on which the work is being performed, uses Xilinx parts to implement most of the glue logic in the PEs (processing elements). These can act as programmable performance monitors at each processor. These tools are being used to evaluate the impact of combining on overall system performance, measure the performance of scientific applications, measure and compare alternative operating system designs, and construct mathematical models of parallel system behavior.

New York University; Zvi M Kedem; *High Performance Parallel Processing: Fault-Tolerant Computing on a Network of Workstations*; (CCR-9411590); \$25,000; 12 months; (Joint support with the Computer Systems Architecture Program - Total Grant \$53,469).

This project deals with both theoretical and prototyping research in parallel processing to harness the power of workstation clusters. The primary techniques will be based on aggressive scheduling, evasive memory and dispersed data management. The formal work expands the previous results of the PIs in asynchronous computing in several.

1. Extend the models to fine-grained programs.
2. Permit the inclusion of architecture specific characteristics into the theoretical framework.
3. Techniques to distinguish inconsistencies in memory states as viewed by different processors resulting from time-outs, interrupts and failures (or indefinite postponement of a computation).
4. Techniques to allow non-deterministic executions.
5. Model input/output operations of fault-tolerant programs.

The experimental part of the project will build a scaled version to execute parallel programs on clusters of workstations. The prototype will address issues related to the eager scheduling of threads, the implementation of evasive memory and data management. The prototype will attempt to evaluate the effectiveness of the environment in terms of efficiency.

This project is conducted jointly with P. Dasgupta of Arizona State University.

Carnegie-Mellon University; Daniel P Siewiorek; *Collaborative Research: Architecture, Design and Implementation of Mobile Computers*; (MIP-9403473 A001); \$10,000.

This is a joint effort between two universities for rapid prototyping of mobile computers. The projects involve teams of students who over the course of a semester design the hardware, software, and packaging of mobile computer systems, and fabricate prototypes by combining standard electronic parts with custom fabricated cases and interconnect harnesses. This project focusses primarily on mobile computing for the diagnosis of telecommunications networks. Mobile computer systems that include stationary and mobile computers, interface devices, and software, are being designed and fabricated for these applications. The project goal is to systematize the prototyping process in several ways: by providing an integrated design tool that can simultaneously represent electronic, thermal, and mechanical constraints; by providing modular template architectures for mobile computers; and by reporting on experience in the co-design of software, mobile hardware, and stationary hardware in a mobile computing system.

University of Washington; Theodore H Kehl; *Self-Timed Logic in Multiprocessors*; (MIP-9101464 A003); \$25,000; 6 months.

This project involves building and measuring two self-timed components inserted into an existing shared-memory multiprocessor computer system. The two components are a multilayered backplane with self-timed arbitration logic and a self-timed memory module. The goals are to demonstrate the ability to increase the number of processors while also doubling memory performance for this system. This project is testing the viability of self-tuning systems (the system is self-tuning in that the operating margins are adjusted based on the actual components used in the system).

University of Washington; Lawrence Snyder, Carl Ebeling; *Chaotic Routing: Study and Implementation*; (MIP-9213469 A002, A003); \$411,586; 12 months.

The chaotic router for multiprocessor systems avoids congestion in message routing by derouting packets chosen at random at congested nodes of a network. The routers can thus adapt to varying message traffic. In this project, the router is being implemented and its performance is being measured within a testbed that approximates a real multiprocessor.

University of Wisconsin; Mark D Hill, James R Larus, David A Wood; *Cooperative Shared Memory and the Wisconsin Wind Tunnel*; (MIP-9225097 A002, A003); \$643,137; 12 months.

The goal of this project is to design hardware and software for scalable shared-address-space computers. Cooperative shared memory, the approach taken in the project, provides a simple design for shared-memory hardware and a programming model that can be used by programmers and compilers to understand an application's communication behavior. Cooperative shared memory uses simple directory hardware together with a set of pragmas for use in applications software. The pragmas allow the applications software to indicate which processors will be using a block of memory: a processor can check out a block when it expects frequent use, check it back in when it is done, and can indicate that it expects to check out a block in the near future. Simple directory hardware can be used to place checked-out-memory locations in the caches of the appropriate processors. Common state transitions in the directory protocol are implemented in hardware, while the less common ones use software traps. Note that the pragmas and the resulting hardware actions affect only the execution speed of a program, not its correctness. A new virtual prototyping approach is being used for evaluating the new architecture. The Wisconsin Wind Tunnel runs parallel shared-memory programs on a parallel message-passing computer and concurrently evaluates the programs execution times on proposed hardware using a distributed simulation. The simulation runs quickly because instructions that make only local memory references are executed directly.

University of Wisconsin; Gurindar S Sohi, James E Smith; *Prototyping Multiscalar Processors*; (MIP-9505853); \$181,670; 12 months; (Joint support with the Computer Systems Architecture Program - Total Grant \$261,670).

This project is evaluating a new architectural paradigm that can extract and exploit the parallelism in sequential code. This new approach uses both software scheduling in the compiler, as in VLIW, and hardware scheduling at run-time, as in superscalar architectures. The compiler segments code into large blocks of instructions that form subgraphs of the control-flow graph, though not necessarily basic blocks. The compiler appends synchronization information to each block that describes which registers must be shared with blocks. The compiler appends synchronization information to each block that describes which registers must be shared with other blocks. The blocks or tasks are passed to separate identical parallel execution units, each of which executes its task sequentially. Tasks are scheduled optimistically, so that some tasks may be executed by a unit when they would not be executed in a sequential system; in such cases, the tasks are

"squashed," which means that their results are not written to memory or registers. At the end of each task, it blocks until the system determines that the task will not be squashed; at that time results are committed. The collection of execution units appears logically to be one unit, with a single register file. Shared register values are passed on an inter-execution-unit network, and tasks that need shared values block until the values are produced.

During the first two years, this research is devoted to compiler design for the new architecture, to comparison with other architectures using simulation, and to conceptual refinement of the architecture. During later years, work is expected to focus on implementation.

Application Specific Computing

International Computer Science Institute; Nelson Morgan; *A System for Connectionist Speech Recognition Research*; (MIP-9311980 A001); \$359,999; 12 months.

This project is constructing a computer optimized toward speech recognition algorithms, and is evaluating speech algorithms on the machine. A fundamental goal of the project is to explore the architectural changes needed for speech processing in future production systems. The new computer is a low-degree multiprocessor, each node of which contains a high-speed general-purpose processor, a multiple-accumulate processor, memory, and a communications interface for the multiprocessor interconnect. The computer will also be capable of being extended to include analog processing or smart sensors. This new machine will provide the performance of supercomputers at a small fraction of their cost on the speech recognition problem, and will contribute to the development of speech recognition systems for everyday use in commodity computers.

processors. The processing elements have been tuned to sequence comparison by incorporating a single cycle add-and-min instruction and a data-path for quickly recording the results of the instruction. Software, I/O, and algorithms have been considered in the design of this architecture; the resulting balanced implementation should result in high performance at low cost.

Colorado State University; Tom Chen, V. Chandrasekar; *An Expandable Column FFT System and Its Multi-Chip Module Implementation*; (MIP-9204319 A001); \$35,000; 12 months.

This project is to explore column Fast Fourier Transform (FFT) architecture in the bit-serial computation and multi-chip module implementation paradigms. The architecture employs flexible routing structures to allow for a wide range of FFT transform lengths and to facilitate defect and fault-tolerance. The object of this project is to build a prototype of the proposed column FFT architecture and demonstrate its performance and expendability on real-time applications (using a 10 cm Dual-Polarized Doppler radar system for real-time meteorological experiments).

University of California-Santa Cruz; Richard Hughey, Kevin Karplus; *Multi-Purpose Parallel Process for Biosequence Analysis*; (MIP-9423985); \$145,934; 12 months.

This project is building an application-specific computer system for biosequence comparison. The architecture is suitable for a wide range of sequence comparison methods, including the Smith-Waterman algorithm, BLAST, profile searches and dictionary methods. In addition, the architecture is being integrated with software for the statistical analysis of sequences using techniques such as hidden Markov models.

The core of the architecture is a linear array of SIMD processing elements, each with a small local memory. A single chip can contain 64 of these elements, so a board with 20 chips will hold 1280

Massachusetts Institute of Technology; Robert C Berwick; *High Performance Computing for Learning*; (IRI-9217041 A003); \$42,000; 12 months; (Joint support with the Knowledge Models and Cognitive Systems Program, the Interactive Systems Program, the Robotics and Machine Intelligence Program, the Linguistics Program, and the Advanced Research Projects Agency - Total Grant \$599,600).

This project has been designed to push the High Performance Computing algorithmic and architectural envelope via a CM-5 and VLSI testbed. It will advance new algorithms and software for a broad class of optimization and learning problems, tested on and directly driving operating system and architectural changes on the CM-5 (working with one of the CM-5's key architects). The learning problems addressed are essentially an entire class of modeling/optimization problems that intersect with nearly all HPCC Grand Challenge Problems.

University of Michigan; Kang G Shin; *Architecture and OS Support for Real-Time Fault-Tolerant Communication*; (MIP-9203895 A003, A004); \$330,791; 12 months.

This is a project to build and experiment with a multiprocessor for real-time applications. The multiprocessor will ultimately consist of 19 nodes arranged in a hexagonal mesh, with each node containing three 68020 processors for computation, a commercial controller for communications control, and a custom chip for communications routing. An initial multiprocessor containing only a few nodes is being developed to allow experimentation that will guide full-scale construction. Experiments on the system consist largely of synthetic benchmark programs produced by a workload generator. The workload generator produces benchmarks that have computation, communication, and deadline characteristics of several classes of applications, but are easier to vary an instrument than real applications would be.

University of Michigan; A. Galip Ulsoy, Yoram Koren, Kang G Shin; *Hierarchical Controller for Real-Time Quality Control in Machining*; (DMI-9313222 A001, A002); \$30,000; 12 months; (Joint support with the Robotics and Machine Intelligence Program, and the Manufacturing Machines and Equipment Program - Total Grant \$145,678).

The award is for the design and development of an experimental, modular, open-architecture machining system controller for prismatic and rotational parts that enhances part quality by an order of magnitude while sustaining high machining rates. The approach would enable integration of underutilized results of the prior work of various researchers as well as generating new results. The hypothesis of the research is that information embedded in the control-loop error signals can be effectively used to improve part quality and machining system performance. A hierarchical controller structure is proposed consisting of compensators at the machine tool servo, process, and part inspection levels. Intelligent integration of hardware and software components of the controller is proposed here by combining expertise in machine tools, machining processes, sensors, control theory, software architecture, and real-time computing.

The hierarchical, open-architecture controller is expected to have a major impact on manufacturing practice and research. Potential benefits for industrial users include reduction of part manufacturing cycle times, reduction of indirect inspection costs, controller improvement costs, and the cost of using sensors. Implementation of future machining research results would also be facilitated by the availability of such a controller.

Princeton University; Andrew Wolfe, Wayne H Wolf, Bede Liu; *An Experimentally-Designed Video Signal Processor Architecture*; (MIP-9408462, A001); \$351,548; 24 months.

This project is designing an architecture and compiler for video signal processing, with collaboration among researchers in computer architecture, compilers, and signal processing. The architecture is a very long instruction word machine dedicated to the applications, which means that a large number of small data paths are used. Address generation may receive hardware support. The compiler differs in several ways from general purpose compilers: optimization is sought even at the cost of high compilation times, and memory usage will be explicitly optimized. Architecture and compiler design are being evaluated by experimentation with application codes, including video compression, image enhancement, and machine vision. The architecture is being simulated on a network of high-performance workstations, while the compiler is being constructed based on the Gnu C compiler. The architecture and compiler will be iteratively improved to result in a high-performance system and a set of design principles.

Cornell University; Herbert B Voelcker; *Massively Parallel Computation for Mechanical Manufacturing and Design*; (MIP-9317620 A001); \$401,329; 12 months.

Solid modeling is a critical enabling technology for mechanical CAD/CAM because it provides geometrically complete representations of parts and products, and enables important manufacturing processes to be modeled. Today's industrial systems operate far below the technology's potential because solid modeling requires enormous computing resources, and because current algorithms and representations cannot handle several important applications.

This project is designing new implementations of ray-casting representations for the solid models used in mechanical CAD, and extending the models to new applications. New implementations for ray-casting engines are being developed, using both custom hardware and software running on massively parallel machines. Using the new hardware, new applications of ray representations are being explored, including the solution of boundary-value problems, computation of medial axis transforms, methods for representing mechanical tolerances, and representation of solid objects.

State University of New York - Stony Brook; Arie E Kaufman; *Scalable Architecture for Real Time Volume Rendering*; (MIP-9527694); \$172,645; 12 months.

This research deals with the development of a new scalable volume visualization architecture and its associated algorithms. The architecture, called Cube-4, exploits parallelism and pipelining to achieve real-time rendering of high-resolution images from volume data. The architecture is based on an algorithm for ray casting of a volume buffer of voxels which is stored as a skewed distributed memory to support conflict-free access to voxel structures. It performs interpolation of sampled points, shading, and compositing to generate the pixel values. Computations are done using limited

communication between processors, so that the architecture is scalable over a wide range of performances and image resolutions. This project encompasses algorithm development, architecture research, and construction of a reduced-resolution prototype.

University of North Carolina; John W Poulton, Henry Fuchs; *Scalable Graphics: From Personal to Supercomputer Visualization Engines*; (MIP-9306208 A005, A006); \$672,105; 12 months; (Joint support with the Advanced Research Projects Agency - Total Grant \$1,409,257).

The object of this project is to build and experiment with a new graphics engine that will eliminate the current limits to scalability in commercial graphics systems. The work centers on a new graphics engine architecture called image composition, which is radically different from the organization of today's commercial systems. In image composition, rendering is distributed over a number of identical processors. Each renderer generates a full-screen image, but for only a fraction of the primitives in the scene. The system then merges these images over a high-speed network to form a single image of all primitives. Since each subimage is independent, and since the images can be merged on a distributed network whose throughput scales linearly with the number of subimages, performance of the entire system can be scaled up arbitrarily by adding more processors.

Carnegie-Mellon University; Susan Finger, Lee Weiss, Daniel P Siewiorek, Andrew P Witkin; *Rapid Design Through Virtual and Physical Prototyping*; (MIP-9420396); 12 months; (Joint support with the Computer Integrated Engineering Program, and the CISE Institutional Infrastructure Program - Total Grant \$499,417).

This project is creating an experimental system using the Internet that will allow students in design courses to use rapid prototyping services. Three participant institutions, Carnegie-Mellon University, the University of California at Berkeley, and Stanford University are participating in the project. Each has developed individual technologies for virtual and physical prototyping, which currently stand alone. Bringing these technologies together will result in exciting new capabilities. Partnerships with industrial partners and a Federal laboratory are also anticipated. This project addresses key issues in prototyping, and is creating a network of interconnected services to support the rapid design, test, and manufacture of mechanical, electro-mechanical, and electronic products.

Carnegie-Mellon University; Roy A Maxion, Andrzej J Strojwas, David L Banks; *Discovering Information in Large,*

***High-Dimensional Databases*; (IRI-9224544 A002); \$10,000; 12 months; (Joint support with the Knowledge Models and Cognitive Systems Program, the Database & Expert Systems Program, the Design, Tools and Test Program, and the Statistics Program - Total Grant \$160,000).**

Discovering functional relationships among high-dimensional data is astonishingly hard. Overcoming the "curse of dimensionality" is a vital problem for any complex manufacturing industry, such as VLSI production, in which hundreds of variables must be precisely controlled in order to achieve high-quality yield. This research addresses both theoretical and practical concerns. On the theoretical end, statisticians have recently proposed a number of compelling new ideas for high-dimensional, nonparametric regression (e.g., ACE, AVAS, LOESS, PPR, MARS, RPR and several other algorithms). These ideas are largely untested, and little is known about their comparative performance in realistic situations. To remedy this, a large-scale simulation experiment is performed that employs statistical design to evaluate the effects of sample size, dimensionality, signal-to-noise ratio, and various kinds of underlying functions on the integrated mean squared error of the fitted model. The results of the study are examined in an analysis of variance, leading to clear conclusions as to the circumstances under which each of the proposed methods is most valuable. On the practical side, this research applies the methodology studied in the simulation experiment to VLSI production data as micro-modeled by the PREDITOR software, which is widely used in industry to calculate from physical principles the actual result of each step in the production of a VLSI circuit wafer.

Carnegie-Mellon University; D. Lansing Taylor, Scott E Fahlman; *High Performance Imaging in Biological Research*; (BIR-9217091 A005, A006); \$100,000; 12 months; (Joint support with the Knowledge Models and Cognitive Systems Program, the Robotics and Machine Intelligence Program, and the Databases, Software Development and Computational Biology Program - Total Grant \$268,000).

This is a project to research and develop an Automated Interactive Microscope (AIM). The AIM will combine the latest technologies in light microscopy and reagent chemistry with advanced techniques for computerized image processing, image analysis, and display, implemented on high-performance parallel computers. This combination will produce an automated, high-speed, interactive tool that will make possible new kinds of basic biological research on living cells and tissues. While one milestone of the research will be to show the proof-of-concept of AIM, the on-going thrust will be continued development as new technologies arise and the involvement of the biological community.

University of Utah; John M Hollerbach, Stephen C Jacobsen, Elaine Cohen; *Rapid Virtual Prototyping of Mechanical Assemblies*; (MIP-9420352); \$408,000; 24 months; (Joint support with the Systems Prototyping and Fabrications Program - Total Grant \$800,000).

This project is building a haptic interface to a mechanical design system. The objective is to demonstrate that part interaction, assembly, and manipulability can be evaluated without physical prototypes. Visual interactions are provided by a CRT screen, while haptic interaction is by means of force reflecting telerobotic master arms, which the user wears as an exoskeleton. Instead of controlling slave arms in a remote environment, the masters affect a virtual world in which the prototype assembly exists. Geometric and functional reasoning determine part contact, interference, and the propagation of forces and torques. Visual rendering techniques display the current state of the assembly. Haptic rendering techniques feed back appropriate forces to the master arms, enhancing the effectiveness of the interface.

Other

Salk Institute for Biological Studies; Terrence J Sejnowski; *Workshop on Neuromorphic Engineering; June 25, 1995 - July 8, 1995; Telluride, Colorado; (IBN-9511637); \$5,000; 12 months; (Joint support with the Interactive Systems Program, the Neuroengineering Program, the Computational Neuroscience Program, the Databases, Software Development and Computational Biology Program - Total Grant \$51,840).*

Recently a new field of engineering has emerged, referred to as neuromorphic engineering, that is based on the design and fabrication of artificial neural systems, such as vision chips, head-eye systems, and roving robots, whose architecture and design principles are based on those of biological nervous systems. A two-week workshop on neuromorphic engineering will be held to bring together young investigators and more

established researchers from academia with their counterparts in industry and national laboratories, working on both neurobiological as well as engineering aspects of sensory systems and sensory-motor integration. Formal lectures will be given, but the primary focus will be hands-on experience with research tools for all participants. The workshop will serve as a bridge between the engineering world of artificial neural systems and the neuroscience community. The workshop will provide an environment for intensive interactions between members of these two communities - merging engineering principles with experimental results from neuroscience. The interaction of these two disciplines should have significant impact both on the development of new technologies (new artificial neural systems) and on our understanding of how the nervous system is designed.

Systems Prototyping and Fabrication

Dr. John Staudhammer, Program Director
(703) 306-1936 jstaudha@nsf.gov

The Program

The Systems Prototyping and Fabrication Program addresses rapid prototyping technologies. The goal is to develop and integrate the tools and technologies needed for rapid and efficient design and fabrication of products, processes and systems. The SPF Program consists of three principal elements. The first (systems prototyping) supports fundamental research on rapid system prototyping methodologies, tools, environments, etc. with particular interest in the informational infrastructure needed for prototyping systems (virtual prototyping

¹). The second (microfabrication) supports research related to advancing the state of the art in the modeling and control aspects of the fabrication (physical prototyping¹). The third element (education) provides assistance to microelectronics education through support of MOSIS and administrative oversight for its involvement in microfabrication for educational institutions.

SYSTEMS PROTOTYPING

Systems Prototyping deals with the information technology for tying together design and fabrication in order to implement quick turnaround design of prototypes. The goal is to create methodologies and define technologies that reduce the time needed to build components and systems. SPF seeks to provide the design methodology and tools as well as infrastructure and services for rapid prototyping. Research is supported on: systems level design tools and environments for virtual prototyping of systems, design frames, design methodology interface problems, specification languages and formats and modeling techniques for packaging technologies.

New initiatives in this program element include research in rapid prototyping for advanced manufacturing. New tools and technologies for virtual prototyping coupled with innovative services and an updated infrastructure that allows distributed rapid prototyping over high speed networks are of particular interest. Projects include new languages for machine and process design, and CAD/CAM integration, as well as projects encompassing modeling, simulation, model validation, and design tools and techniques. Research dealing with the application of useful VLSI design paradigms to SFF (Solid Free-form Fabrication) and MEMS (Micro-Electro-Mechanical Systems) are also addressed.

MICROFABRICATION

This program element supports basic research needed to understand, model and control microfabrication processes, including nanotechnology and biochips, pattern definition and transfer, and modeling, simulation and automation (for computer integrated manufacture of VLSI components and systems). The emphasis is on research focused on modeling and control aspects of physical prototyping. Topics of interest include architectures for manufacturing, simulation and real time control; disciplines for using semiconductor manufacturing equipment and processes; test structures, sensors and instrumentation for process monitoring; modeling and simulation at the process, device, circuit and system levels; integration of CAD, CAM and CAT methodologies and the application of these methodologies to SFF, MEMS and other mechanical and electromechanical prototyping technologies.

EDUCATION

This program element consists of two components. The first is MOSIS (MOS Implementation Service) which serves as a broker providing access to the semiconductor foundry industry. This program also deals with new technology issues at the undergraduate level through sponsorship of educationally oriented conferences and workshops, funding of innovative technology developments that significantly impact educational approaches to system prototyping (such as FPGA design frames), distribution of preliminary versions of innovative research and educational materials, and encouragement for the upgrading of microelectronics-related subject matter, curriculum, laboratories, and faculty.

¹ NEW PARADIGMS FOR MANUFACTURING, NSF Workshop Report NSF-94-123, May 2-4, 1994, Arlington VA.

Initiatives and Opportunities

The Systems Prototyping and Fabrication Program encourages research leading to the development of new technologies for rapid prototyping and provides the research and educational communities with access to these new technologies.

While MOSIS has provided and will continue to provide a valuable service to the university education and research community, other methods of virtual and physical implementation are needed. Methodologies for prototyping, packaging, testing and manufacturing must be integrated with and closely coupled to systems and circuit design technologies. Requirements for higher performance and improved reliability with smaller size, lower cost and lower power also dictate such integration. Designers must work with more of a systems outlook, and have a more comprehensive design experience at a higher level of system implementation.

New technologies (mini-fab production lines, Field Programmable Gate Arrays (FPGA), Multi-Chip Modules (MCM), optical interconnect, micro-sensors, biochips, etc.), and new methodologies (fast prototyping, top-down design, powerful CAD tools, design libraries, etc.) when coupled with innovative services and an updated infrastructure have the potential to meet this need.

Listed below are some of the key issues related to the SPF activities;

- * Overcoming performance limitations increasingly due to packaging shortcomings.
- * Reducing the time and cost for prototype fabrication with new methodologies (virtual prototyping, for example), tools, equipment, and services.
- * Exploiting new technologies (field programmable gate arrays, multichip modules, etc.) and new methodologies for rapid physical prototyping in research and education.
- * Applying relevant VLSI design paradigms to solid free-form fabrication (SFF), micro-electromechanical systems (MEMS) and other mechanical technologies.
- * Simplifying, automating and speeding up access to microfabrication processes.
- * Making new package and multi-chip module approaches available to the academic community.
- * Defining new methods for functional and physical partitioning across and within package levels.
- * Encouraging a closer relationship between CAD tool designers and those doing fabrication, packaging and prototyping, in areas such as requirements, integration, and evaluation of system performance.
- * Developing, integrating and improving distributed access to virtual and physical prototyping techniques.
- * Developing design frames that allow rapid hardware simulation of application-specific systems.
- * Exposing students to a system-level design experience, with practice in optimizing the selection among design alternatives and with exposure to requirements for design verification.
- * Innovative use of curricular materials, compression of topics, curriculum updating, and use of higher levels of abstraction.

Awards

Prototyping

Microelectronic

California Institute of Technology; Michelle Effros; CAREER: Code Clustering for Universal Image Coding and Other Implications; (MIP-9501977); \$32,169; 36 months; (Joint support with the Circuits and Signal Processing Program, the Microelectronic Systems Architecture Program, and the Design, Tools and Test Program - Total Grant \$134,628).

This project seeks to develop reasonable complexity, source-independent coding algorithms, which are crucial to the design of robust systems for image coding and mobile communications. In these applications the statistics of the source and channel in operation are typically unknown a priori, and the performance of the coding strategy employed is sensitive to those unknown characteristics.

The two-stage approach developed in the source coding literature is employed. The literature demonstrates that in general one should quantize the space of possible codes. Some of the rate should be spent on describing which code, in a family of codes, should be used on the source in operation. Specific projects include the development of a universal DCT code compatible with JPEG and MPEG image and video standards, a universal KLT code, a universal wavelet packet code, and a universal channel code.

The main objectives of the education plan are to develop and maintain an exciting atmosphere for active learning for undergraduate and graduate students through innovative programs that encourage the maximum possible exchange between students, faculty, and individuals from local industry.

University of California-Santa Barbara; Malgorzata Marek-Sadowska; Research on Layout and Logic Design; (MIP-9419119); \$40,000; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$140,000).

This research is on layout driven synthesis, i.e. the intersection of logic synthesis and physical design. The focus is on restructuring logic networks in synthesized digital systems. Five topics, which meet the goals of improving routing efficiency or power consumption, are being investigated. These are:

1. Incremental logic resynthesis to control wiring;
2. Coupling wiring with logic restructuring and finding optimizations to eliminate wiring overflows;

3. Use of generalized Reed-Muller forms to analyze logic as an aid to;
 - a. designing cell libraries,
 - b. technology mapping,
 - c. developing new multi level optimization techniques,
 - d. designing networks of provably good testability;
4. Develop new methods for power optimization, at the technology independent and technology dependent levels in logic synthesis; and
5. Find better routing tools to handle power constraints.

University of Florida; Mark E Law; PFF: A Multidisciplinary Approach to IC Process Modeling Using the SUPREM-IV Modeling Tool; (MIP-9253735 A003); \$100,000; 12 months.

This multidisciplinary research focuses on the development of silicon models for point defect behavior, which are vital to understanding dopant diffusion. Models are being developed and parameterized for the effect of silicidation and stress on point defect kinetics. These models are then implemented in SUPREM-IV, a standard integrated circuit process modeling tool that utilizes advanced finite element techniques.

University of Maryland; K. J. Ray Liu; NYI: High Performance Computing for Signal Processing; (MIP-9457397 A001); \$31,250; 12 months; (Joint support with the Circuits and Signal Processing Program - Total Grant \$62,500).

There are three major architectural models used in high-performance signal/image processing:

1. VLSI- signal processing - high- throughput VLSI architectures for low-cost application-specific implementations used in applications such as communication systems, speech, video/HDTV, and radar;
2. parallel signal processing on massively parallel computers - parallel algorithms for complex signal/imaging systems used in computer vision, medical imaging, and the processing of vast amounts of data in deep space exploration;
3. distributed signal processing on high-speed networks - used in applications such as document image processing, multimedia, automatic signal processing in manufacturing, and medical signal/image processing.

This research will focus on the development of efficient algorithms and architectures for each architectural model, and in comparative studies of the advantages and disadvantages of these different computing schemes. The goal is to investigate which

signal/image processing problems can be carried out optimally under different computing and communication schemes.

University of Maryland; Linda Milor; CAREER: A Statistical Modeling Methodology for Submicron MOS Devices and Circuits; (MIP-9501912); \$135,000; 36 months.

The performance of an integrated circuit depends on the designers' choices and the process engineers' decisions during the manufacture of the device. Typically the choices and decisions are independent of each other with little awareness of the interplay of the two engineering domains.

This research develops methods for raising the designers' awareness of the effects of circuit implementation parameters and to sensitize process engineers to the impact of processing decisions on the performance of the final circuit. A statistical model is developed to correlate device and process parameters, primarily to establish circuit performance variations on spatial and process dependencies in replicated circuits. Circuit simulators are used to solve the resultant

Michigan Technological University; Ashok K Goel, Esther T Ososanya; Experimental Validation of Interconnection and Transistor Delay Models for the GaAs-Based Integrated Circuits; (MIP-9223989 A003); \$10,000.

During the last few years, GaAs technology has emerged rapidly from basic research to device and circuit development. It is crucial to know the expected propagation delays in an integrated circuit before it is fabricated. To meet this objective, numerical models have been developed that address crosstalk and propagation delays in the parallel and crossing VLSI multilevel interconnections as well as for the transverse delays in the GaAs MESFETs and GaAs/AlGaAs MODFETs. In addition to determining the crosstalk and propagation delays, the models can be utilized to achieve the optimization of the device and interconnection dimensions and other parameters for minimum crosstalk and delays. Validation of these numerical models by comparison of the modeling results with actual experimental observations is critical if they are to be incorporated into GaAs CAD tools. This research effort focuses on the following set of objectives:

1. design and fabrication of several GaAs-based logic circuits to retain the ability to alter the various design parameters;
2. application of the interconnection and the GaAs MESFET delay models recently developed for the determination of propagation delays in these GaAs-based logic circuits;
3. experimental measurements of propagation delays in these circuits and comparison with developed

- delay models;
4. modification of the interconnection and transistor delay models, as required; and
5. experimental validation of the final models.

State University of New York - Binghamton; Jiayuan Fang; NYI: Analysis and Modeling of High-Speed Interconnects in Electronics Packaging; (MIP-9357561 A002); \$61,984; 12 months.

This research is concerned with the analysis and modeling of electrical performance of high-speed interconnects in electronics packaging. The finite-difference time-domain (FDTD) method, which is a full-wave solution of Maxwell's equations in three dimensions, is used to simulate signal propagation through interconnects. Topics pursued are:

1. Development of a computational scheme for conformed finite-difference grid to model complex-shape interconnects. The objective of this scheme is to enhance the resolution and accuracy of numerical solutions while maintaining the computation efficiency associated with the regular rectangular finite-difference grid.
2. Analysis and modeling of electrical properties of interconnection discontinuities in electronics packaging. Issues involved in this topic include: modeling of electrical characteristics of interconnection discontinuities over the frequency range from dc to tens of gigahertz; evaluation of impacts of parasitics associated with interconnection discontinuities on the propagation of high-clock rate signals; and development of design guidelines for typical interconnection discontinuities in high performance electronics packaging.

North Carolina State University; Paul D Franzone; NYI: Advanced Interconnect and Display Approaches; (MIP-9357574 A002); \$62,500; 12 months.

The primary focus of this work is to resolve issues dealing with the design and implementation of high bandwidth reconfigurable interconnect systems based on Micro Electro-Mechanical Systems (MEMS) which are commonly referred to as Micromachines. Different guided-wave optical and holographic free-space optical switch elements are being implemented and compared in terms of bandwidth, switch reconfiguration rate, and range of programmability.

Application to data switching, and programmable interconnect devices for rapid prototyping are being addressed with attention to both technological and system-wide performance/cost design issues. Comparisons are made with conventional technologies. Also being investigated is the application of some of the optical MEMS elements to advanced image projection.

University of North Carolina - Charlotte; Dian Zhou; NSF Young Investigator: Performance-Driven VLSI Designs; (MIP-9457402 A001); \$31,250; 12 months; (Joint support with the Design, Tools and Test Program - Total Grant \$62,500).

Three research issues in high-performance VLSI system design are being addressed:

1. how to relate the system performance function, characterized by electrical parameters, to the geometrical parameters of the VLSI physical design;
2. how to model performance driven VLSI physical designs based on given technology and computational capability; and
3. how to characterize the fundamental computational aspects of modeled problems and develop effective algorithms for solving them.

A distributed RLC circuit model for interconnects is being designed. It considers:

1. non-monotone circuit response;
2. coupling effects among the signal lines, and
3. low energy consumption.

Efficient computation methods that solve time-varying Maxwell equations using the adaptive wavelet collocation method (AWCM) are being devised. The algorithms and methods are being included in a CAD system.

Brown University; Harvey F Silverman; *A Large-Scale, Intelligent Three-Dimensional Microphone-Array Sound-Capture System*; (MIP-9314625 A002, A003, A004); \$87,940; 24 months; (Joint support with the Experimental Systems Program, the Circuits and Signal Processing Program, the Microelectronic Systems Architecture Program, and the Design, Tools and Test Program - Total Grant \$674,601).

This is a joint project between Brown and Rutgers to build a large microphone array with associated processors for beam forming. Applications are direction finding, echo cancellation, and speaker differentiation in teleconference systems, multimedia educational systems, and large conference centers. Groups of microphones share microphone modules that perform A/D conversion and low-level processing. The microphone modules are linked over a high speed serial network (possibly optical) to a multiprocessors signal processing system for the higher level processing. The resulting system is being evaluated in an experimental teleconferencing facility.

University of Wisconsin; B. Ross Barmish, Krishna Shenai; *Robust VLSI*; (MIP-9424580); \$25,000; 12 months; (Joint support with the Systems Theory Program - Total Grant \$50,000).

The focus of this research is on the design and manufacture of electronic circuits at the submicron range for logic, analog, and mixed-signal applications. This new approach to robust design uses an interval bound description rather than statistical description of variable parameters associated with the manufacturing process. Evaluation of performance utilizes advanced mixed-mode circuit simulators in which carrier dynamics within semiconductor devices and interconnects are modeled at the application level. The focus is on the design of algorithms with the long term goal of developing tools to aid in the analysis and design of complex VLSI circuits. Initial examples being considered are a CMOS differential amplifier and a CMOS inverter.

MEMS, SFF, MCM

Stanford University; Martin Fischer, Krishna C Saraswat, Raymond E Levitt; *Product, Process and Organization Prototyping for Concurrent Engineering*; (MIP-9420398); \$150,001; 24 months; (Joint support with the Information Technology and Organizations Program - Total Grant \$300,001).

This project conducts enabling research on virtual prototyping tools for collaborative design, construction, and startup of semiconductor manufacturing facilities.

This extends and integrates ongoing research within the Center for Integrated Facilities Engineering and the Center for Integrated Systems at Stanford. The focus is not on design of actual projects or processes, but on the implementation and startup of the factory building and support systems of the manufacturing enterprise.

Research is being performed in collaboration with an industrial partner who is providing specifications and is guiding testing of theory and software. The resulting tools will radically shorten the time to bring new semiconductor facilities on-line.

University of Colorado; Yung-Cheng Lee; *PYI: Multichip Module Design for Manufacturing*; (MIP-9058409 A008); \$47,500.

This research addresses some key requirements for the design for manufacture of very small supercomputers used in intelligent machines such as portable robots. This multidisciplinary effort is centered on developing a compact rapid prototyping and manufacturing center. Microscale laser lithography, flip-chip soldering and robot controlled pick-and-place techniques are being used. Simulation studies validate the design before prototyping. Research is concentrating on the self-aligning mechanism with an emphasis on the reliability of solder joints, fuzzy logic modeling with focus on the process modeling and optimization, and further improvement of thermosonic bonding technology.

Massachusetts Institute of Technology; Emanuel M Sachs, Duane S Boning, Michael J Cima, Subra Suresh, David C Gossard; *Design Automation for Solid Freeform Fabrication*; (MIP-9420365); \$360,000; 24 months.

This project is creating a design paradigm for solid free-form fabrication, based on that for VLSI. Elements of the VLSI paradigm to be translated to the mechanical domain include one-directional information flow, the clear delineation of levels of abstraction, formulation of design rules and design libraries, and the development of intermediate representations for storage

and transmission of various levels of abstraction of a design. Several types of structures that are ill-suited to present design methods are the motivation for new languages, libraries, and tools. These include surface textures, functionally gradient materials, and thermally and chemically active materials. All of these structures require instantiation of cells, possibly parameterized, which is not well supported by current languages such as STL.

Michigan State University; Jon Sticklen, James K McDowell, Martin C Hawley; *Virtual Prototyping for Polymer Composite Assemblies*; (MIP-9420351); \$525,016; 36 months.

The project develops a virtual prototyping facility for the design of mechanical assemblies containing polymer composites. Design approaches are investigated for lowering the cost of manufacturing, especially methods which replace mechanical subassemblies by a unitary composite part. A suite of automated design tools is developed for virtual prototyping of assemblies made of polymer composite materials AND making these tools accessible on the Internet. Industrial relevance is insured by direct participation of a midsized company specializing in advanced polymer composites design and fabrication.

Carnegie-Mellon University; Jonathan Cagan, Cristina H Amon, Rob A Rutenbar; *Virtual Rapid Prototyping of 3-Dimensional Electro-Mechanical Layouts*; (MIP-9420372); \$469,779; 24 months.

This project is developing techniques and tools for synthesizing three-dimensional layouts of electromechanical systems. Research in this project is exploring new combinatorial algorithms for the arrangement of mechanical components that will satisfy a large number of conflicting design constraints, including volume, shape, weight, thermal interactions, structural support, and accessibility for repair. These algorithms are being incorporated into tools for the design of wearable computers, to be fabricated under the auspices of other NSF-supported research at Carnegie Mellon University. To improve the educational aspects of this project and the connection to the realities of product design, each graduate student on the project will work as a summer intern for industrial supporters of the project in the thermal analysis and project layout groups.

University of Pennsylvania; Ruzena K Bajcsy, Dimitri Metaxas, Vijay Kumar, Daniel K Bogen; *Rapid Prototyping of Rehabilitation Aids for the Physically Disabled*; (MIP-9420397 A001); \$10,000.

The methods of rapid prototyping are ideally suited to rehabilitation devices. Because each person requires unique performance and function in a rehabilitation device, devices specific

to each person must be rapidly designed and produced. This project is investigating a completely integrated approach to the design and prototyping of passive mechanical rehabilitation devices. The approach involves: the quantitative assessment of the form and performance of human limbs; the design of the assistive device; evaluation of the device using virtual prototyping; feedback from the consumer and therapist; actual prototyping of the device; evaluation of the function and performance of the device; and redesign based on performance. The contributions of the product include: the development of new computer-based tools for the assessment of human performance; a manufacturing technique for a new class of hyperelastic materials; the integration of tools into a rapid prototyping system for rehabilitation devices; and development of mechanisms for systematic evaluation of the final product.

University of Utah; John M Hollerbach, Stephen C Jacobsen, Elaine Cohen; *Rapid Virtual Prototyping of Mechanical Assemblies*; (MIP-9420352); \$392,000; 24 months; (Joint support with the Experimental Systems Program - Total Grant \$800,000).

This project is building a haptic interface to a mechanical design system. The objective is to demonstrate that part interaction, assembly, and manipulability can be evaluated without physical prototypes. Visual interactions are provided by a CRT screen, while haptic interaction is by means of force reflecting telerobotic master arms, which the user wears as an exoskeleton. Instead of controlling slave arms in a remote environment, the masters affect a virtual world in which the prototype assembly exists. Geometric and functional reasoning determine part contact, interference, and the propagation of forces and torques. Visual rendering techniques display the current state of the assembly. Haptic rendering techniques feed back appropriate forces to the master arms, enhancing the effectiveness of the interface.

West Virginia University; Lawrence A Hornak; *NYI: Cointegrated Polymer Waveguide Optical Interconnections for Wafer-Level MCM Systems*; (MIP-9257101 A003); \$62,500; 12 months.

Motivated by the need for robust polymers optically superior to polyamides yet suitable for cointegration of optical interconnection waveguides directly with the active CMOS interconnection substrate of advanced multi-chip modules (MCMs), the research seeks to fabricate the first optical waveguides with polyphenylsilsesquioxane (PPSQ), a spin castable, low temperature processed silicon ladder polymer used as a thin and thick film dielectric for microelectronics. This material potentially offers the low loss obtainable with less stable optical polymers while offering thermal stability and process latitude (patterning, wet, dry etching) exceeding that of polyamide. This research assesses the suitability of PPSQ together with planarizing and superstrate polymer layers for providing a multilayer system supporting fabrication of high density waveguide arrays directly over the Silicon devices of emerging active MCM substrate.

Education

Film Arts Foundation; Ruth Carranza; *Silicon Run I.2 - An Educational Film about Integrated Circuits*; (DUE-9455761); \$12,500; 24 months; (Joint support with the Course and Curriculum Program - Total Grant \$50,000).

Integrated circuits, and the computers they drive, introduced a new era in electronics. Each year faster, more powerful computers are produced increasing the demand for a skilled work force and an educated public. Increasing an awareness about manufacturing processes in the semiconductor and computer industries is an educational challenge - one which is hindered by the loss of visual imagery due to the inaccessibility of industrial sites to students, educators and the public.

The SILICON RUN two-part series addresses this educational challenge by capturing the rich images of manufacturing on film/video to provide a clear

overview of semiconductor and computer manufacturing. SILICON RUN I, produced in 1986, explores crystal growth, the effects of doping, CMOS transistors, and the design and fabrication of IC wafers.

The recently completed SILICON RUN II continues with testing, packaging, printed circuit board and system assembly, and multi-chip module technology. The series is being used for education and training; thousands of students in universities, colleges and technical schools are viewing these films/videos every year.

The many changes in technology since the completion of SILICON RUN I make it necessary to update the first film with the production of SILICON RUN I.2. The goal of this project is to re-film the industrial manufacturing footage so the series can continue providing students with viable information about today's manufacturing. Responses from recipients' surveys overwhelmingly

show that the series is an asset to undergraduate education.

Loyola College in Maryland; Michael J DeHaemer; *International Research and Development: What the U. S. Can Learn Abroad*; (ENG-9416970, A001, A002); \$12,123; 24 months; (Joint support with the Interactive Systems Program, the Advanced Research Projects Agency, the Department of Commerce, the Manufacturing Machines and Equipment Program, the Quantum Electronics Program, the Plasmas and Electromagnetics Program, the Special Studies and Analyses Program, the Department of Energy, and the Department of the Navy - Total Grant \$988,945).

This award implements a plan for the Japanese Technology Evaluation Center and the World Technology Evaluation Center (JTEC/WTEC) for conducting twelve studies of foreign R&D to establish benchmarks to help focus US R&D. The research will be done by expert panels that will visit foreign laboratories, and present their findings in workshops and in technical reports. Three methods will be used as appropriate:

1. the standard approach that JTEC/WTEC has used in previous grants, and two new methods that will provide greater flexibility;
2. mission-directed studies on behalf of government agencies that must satisfy a limited set of objectives and meet a short time constraint; and
3. state-of-the art reviews with subjects and delegations proposed by the scientific community.

Carnegie-Mellon University; Daniel P Siewiorek; *Workshop on Solid Freeform Fabrication, May 24-26, 1995, Pittsburgh, Pennsylvania*; (MIP-9522091); \$26,500; 6 months.

A workshop was held involving about a dozen leading researchers from academia and industry met to discuss:

1. The type of part description for use in SFF definition;
2. the nature of design tools to support SFF fabrication; and,
3. the formats for data interchange in solid free-form fabrication.

Infrastructure

Advanced Research Projects Agency; Robert F. Lucas; *NSF/ARPA Agreement for Use of ARPA VLSI Implementations*; (MIP-9419682 A001); \$350,000; 12 months; (Joint support with the Microelectromechanical Research Program, the Manufacturing Machines and Equipment Program, and the USEME Course and Curriculum Program - Total Grant \$800,000).

advanced sensor and micro-mechanical devices); and rapid prototyping methodologies, tools, and services needed for complete systems.

A 1994 Memorandum of Understanding (MOU) between the National Science Foundation (NSF) and the Advanced Research Projects Agency (ARPA) extended a three-year joint program supporting VLSI (Very Large Scale Integration Fabrication) by MOSIS (Metal Oxide Semiconductor Implementation Service) for qualifying educational institutions. The continuation of the MOU expands the original program to accelerate critical capabilities for Microsystems Design and Prototyping in U.S. educational institutions. This includes expanding the original services and technologies available to schools authorized to use MOSIS as they become generally available and cost-effective (e.g., semi-custom and gallium arsenide chip fabrication), stressing VLSI education, especially undergraduate education needed for designing future electronic system; exploring new fabrication services designed specifically for the research and education community's desire for cost effective experimentation of state-of-the-art technologies (e.g., design of

Index of PYI, NYI, CAREER and PFF Investigators

A		N	
Agarwal, Anant	PYI	Nguyen, Truong	CAREER
	58	Nowick, Steven M.	CAREER
			7
B		O	
Baraniuk, Richard G.	NYI	Orchard, Michael T.	NYI
Beerel, Peter A.	CAREEREER	Ortega, Antonio	CAREER
Bresler, Yoram	PYI		43
	44		
C		P	
Chan, Wai-Yip G.	CAREER	Panda, Dhabaleswar K.	CAREER
Chandrakasan, Anantha P.	CAREER	Parhi, Keshab K.	NYI
Chatterjee, Abhijit	CAREER	Pedram, Massoud	NYI
Chiueh, Tzi-Cker	CAREER	Pileggi, Larry T.	CAREER
	31, 44	Prince, Jerry L.	PFF
	6		45
	12		
	57		
D		R	
Douglas, Scott C.	CAREER	Ramachandran, Umakishore	PYI
	41	Riskin, Eve A.	NYI
E		Roy, Kaushik	CAREER
Effros, Michelle	CAREER		6, 24
	15, 30, 42, 69		
F		S	
Fang, Jiayuan	NYI	Sapatnekar, Sachin	CAREER
Farrens, Matthew	NYI	Schulz, Timothy J.	CAREER
Fiez, Terri S.	NYI	Sha, Edwin H.	CAREER
Franzon, Paul D.	NYI	Siu, Kai-Yeung	NYI
	70	Stonick, Virginia L.	PYI
	21		41
	39		
	71		
G		T	
Green, Michael M.	NYI	Torrellas, Josep	NYI
Gupta, Rajesh K.	CAREER		58
Gupta, Sandeep K.	CAREER		
	53		
	5		
	12		
H		U	
Harris, John G.	CAREER	Uehara, Gregory T.	CAREER
Hornak, Lawrence A.	NYI		37
	37		
	73		
K		V	
Kaeli, David R.	CAREER	Vaidya, Nitin	CAREER
Konstantinidou, Smaragda	CAREER	Varma, Anujan	NYI
	32	Vinnakota, Bapiraju	CAREER
	25		13
L		W	
Law, Mark E.	PFF	Wornell, Gregory W.	CAREER
Lee, Yung-Cheng	PYI		45
Li, Jian			
	NYI		
	49		
Liu, K. J. Ray	NYI		
	52, 69		
M		X	
Milor, Linda	CAREER	Xu, Guanghan	CAREER
	70		51
		Z	
		Zegura, Ellen W.	CAREER
		Zhou, Dian	NYI
			8, 71
CAREER	Faculty Early Career Development		
NYI	NSF Young Investigator		
PFF	Presidential Faculty Fellow		
PYI	Presidential Young Investigator		

Index of Principal Investigators

A			
Agarwal, Anant	25,58,59	Cohen, Elaine	64,73
Amon, Cristina H.	72	Cong, Jason	3
Andreou, Andreas G.	32	Cyre, Walling R.	10,30
Armstrong, James R.	10,30		
Aylor, James H.	58	D	
B		Danzig, Peter B.	57
Baer, Jean-Loup	34	Davidson, Jack W.	58
Bajcsy, Ruzena K.	73	Davis, Richard A.	40
Bamberger, Roberto H.	54	DeHaemer, Michael J.	74
Banerjee, Prithviraj	5	DeMicheli, Giovanni	3,58
Banks, David L.	16,64	Devadas, Srinivas	6
Baraniuk, Richard G.	41	Djuric, Petar M.	49
Barmish, B. Ross	71	Donoho, David L.	48
Barnwell, Thomas P.	43	Douglas, Scott C.	41
Basu, Sankar	46	Dubois, Michel	57
Beerel, Peter A.	4		
Berwick, Robert C.	62	E	
Bhuyan, Laxmi N.	28	Ebeling, Carl	60
Blum, Rick S.	50	Effros, Michelle	15,30,42,69
Bogen, Daniel K.	73	Ercegovic, Milos D.	21
Boning, Duane S.	72	Etter, Delores M.	40
Bose, Bella	33		
Bresler, Yoram	44	F	
Brewer, Forrest	3	Fahlman, Scott E.	64
Brockwell, Peter J.	40	Fan, Jianqing	50
Brown, Frank M.	16	Fang, Jiayuan	70
Brunvand, Erik L.	10,29	Farrens, Matthew	21
Bryant, Randal E.	16	Felten, Edward W.	59
Burns, Steven M.	10	Ferguson, Frankie J.	12
C		Fiez, Terri S.	39
Cagan, Jonathan	72	Finger, Susan	63
Carley, L. Richard	57	Fischer, Martin	72
Carranza, Ruth	73	Flynn, Michael J.	58
Chan, Wai-Yip G.	31,44	Fortes, Jose A.	24
Chandrakasan, Anantha P.	6	Franzon, Paul D.	71
Chandrasekar, V.	62	Friedman, Eby G.	8
Chao, Liang-Fang	6	Fuchs, Henry	63
Chatterjee, Abhijit	12		
Chen, Tom	62	G	
Cheng, Chung-Kuan	3,21	Gardner, William A.	48
Cheng, Kwang-Ting	11,22	Geman, Stuart A.	46
Chien, Andrew A.	23	Ghose, Kanad	26
Chiueh, Tzi-Cker	57	Giannakis, Georgios B.	52
Christie, Phillip	16,22	Gidas, Basilis	46
Cima, Michael J.	72	Goel, Ashok K.	70
Clark, Douglas W.	59	Gopalakrishnan, Ganesh	10,29
Clarke, Edmund M.	9	Gossard, David C.	72
Clarkson, Peter	41	Gottlieb, Allan	59
Clements, Mark A.	43	Gray, Paul R.	37
		Gray, Robert M.	

Green, Michael M.	8,53	Lilja, David J.	26
Grenander, Ulf	46	Lipton, Richard J.	59
Gruss, Andrew	57	Liu, Bede	62
Gupta, Rajesh K.	5,23	Liu, C. L.	5
Gupta, Sandeep K.	12	Liu, K. J. Ray	52,69
		Louri, Ahmed	21,30
		Lu, Shih-Lien	27
		Lucas, Robert F.	74
H		M	
Hachtel, Gary	4	Malik, Sharad	7
Harjani, Ramesh	13	Maragos, Petros	40
Harris, John G.	37	Marek-Sadowska, Malgorzata	4,69
Hawley, Martin C.	72	Marron, James S.	50
Hayes, John P.	13,25	Martonosi, Margaret R.	59
Hill, Mark D.	60	Masaki, Ichiro	57
Hollaar, Lee A.	58	Maxion, Roy A.	16,64
Hollerbach, John M.	64,73	McClellan, James H.	43
Hornak, Lawrence A.	73	McClure, Donald E.	46
Hughey, Richard	61	McCluskey, Edward J.	11
		McDowell, James K.	72
		McGrath, S. J.	43
J		Mendel, Jerry M.	39
Jacobsen, Stephen C.	64,73	Menon, Premachandran	7
Jha, Niraj K.	14,26	Metaxas, Dimitri	73
Johnson, C. Richard	41	Meyer, Robert G.	37
Johnson, Don H.	54	Michalopoulou, Zoi-Heleni	32
Johnson, Steven D.	5	Milor, Linda	70
Johnstone, Iain M.	48	Morgan, Nelson	61
K		N	
Kaeli, David R.	32	Nair, Sukumaran	9
Kahng, Andrew B.	3	Nguyen, Truong	48
Kanade, Takeo	57	Nowick, Steven M.	7
Karplus, Kevin	61		
Kassam, Saleem A.	50	O	
Kaufman, Arie E.	63	Olshen, Richard A.	42
Kedem, Zvi M.	60	Orchard, Michael T.	44
Kehl, Theodore H.	60	Ortega, Antonio	43
Kenney, John G.	38	Ososanya, Esther T.	70
Kime, Charles R.	14		
Kogge, Peter M.	24	P	
Konstantinidou, Smaragda	25	Panda, Dhabaleswar K.	27
Koren, Yoram	62	Parhi, Keshab K.	53
Kumar, Vijay	73	Paulraj, Arogyaswami J.	52
Kunz, Wolfgang	14	Pease, Fabian W.	58
Kwong, Dim-Lee	51	Pedram, Massoud	4,22,57
		Pileggi, Larry T.	14
		Pineda, Fernando	32
L		Pinkham, Roger S.	46
Larrabee, Tracy	12	Pomeranz, Irith	13
Larus, James R.	60	Pomerleau, Dean	57
Law, Mark E.	69	Poulton, John W.	63
Lee, Yung-Cheng	72	Pradhan, Dhiraj K.	
Leviton, Steven P.	9,28		
Levitt, Raymond E.	72		
Li, Jian	49		
Li, Kai	59		
Lightner, Michael	4,15		

14Prince, Jerry L. 45

R

Ramachandran, Umakishore 22
Ramanathan, Parameswaran 34
Ranganathan, N. 31
Reddy, Sudhakar M. 13
Riloff, Ellen M. 58
Riskin, Eve A. 47
Robins, Gabriel 10
Rose, Christopher 32
Rosenblatt, Murray 40
Roy, Kaushik 6,24
Rutenbar, Rob A. 72
Ruymgaart, Frits H. 47

S

Saavedra, Rafael H. 57
Sachs, Emanuel M. 72
Sapatnekar, Sachin 25
Saraswat, Krishna C. 72
Sarrafzadeh, Majid 5
Sauer, Ken D. 45
Schafer, Ronald W. 43
Schreier, Richard 38
Schulz, Timothy J. 46
Sechen, Carl 11
Sejnowski, Terrence J. 65
Sha, Edwin H. 6
Shen, John P. 27
Shenai, Krishna 71
Shin, Kang G. 62
Siewiorek, Daniel P. 60,63,74
Silverman, Harvey F. 17,33,53,57,71
Siu, Kai-Yeung 31,52
Smith, James E. 61
Snyder, Lawrence 60
Sohi, Gurindar S. 61
Somenzi, Fabio 4
Song, Bang-Sup 38
Stewart, Alan L. 46
Sticklen, Jon 72
Stonick, Virginia L. 41
Strojwas, Andrzej J. 16,64
Suresh, Subra 72
Swindlehurst, A. L. 51

T

Taylor, D. Lansing 64
Torrellas, Josep 58
Tragoudas, Spyros 12
Trajkovic, Ljiljana 15
Tsatsanis, Michail K. 52
Tugnait, Jitendra K. 39

U

Uehara, Gregory T. 37
Ulsoy, A. Galip 62

V

Vaidya, Nitin 29,33
Varma, Anujan 22
Vetterli, Martin 43
Vinnakota, Bapiraju 13,14
Voelcker, Herbert B. 63

W

Wah, Benjamin W. 23
Walker, Duncan M. H. 14
Walker, Robert A. 8
Walnut, David F. 47
Wei, Belle 30
Weiss, Lee 63
Wey, Chin-Long 13
Windley, Phillip J. 10
Witkin, Andrew P. 63
Wolf, Wayne H. 7,26,62
Wolfe, Andrew 62
Wong, S. Simon 58
Wood, David A. 60
Wooley, Bruce A. 58
Wornell, Gregory W. 45
Wulf, William A. 58
Wyatt, John L. 57

X

Xu, Guanghan 51

Y

Yang, Andrew T. 11
Yang, Qing 28
Yang, Woodward 38
Yang, Yuanyuan 29
Yates, Roy D. 32

Z

Zegura, Ellen W. 23
Zemanian, Armen H. 16
Zhou, Dian 8, 71
Zoltowski, Michael D. 49

Index of Institutions

Advanced Research Projects Agency	74	Stevens Institute of Technology	46
Auburn University	39	Texas A & M University	14, 28, 29, 33
Brigham Young University	10, 51	Texas Technological University	47
Brown University	17, 33, 46, 53, 57, 71	University of Arizona	21, 30
California Institute of Technology	15, 30, 42, 69	University of California-Berkeley	15, 37, 43
Carnegie-Mellon University	9, 16, 27, 41, 57, 60	University of California-Davis	21, 48
.....	60, 63, 64, 72, 74	University of California-Irvine	31, 52
Colorado State University	40, 62	University of California-Los Angeles	3, 21
Columbia University	7	University of California-San Diego	3, 21
Cornell University	41, 63	University of California-Santa Barbara	3, 4, 11, 22, 69
Film Arts Foundation	73	University of California-Santa Cruz	12, 22, 61
George Mason University	47	University of Colorado	4, 15, 40, 72
Georgia Institute of Technology	12, 22, 23, 40, 43	University of Delaware	16, 22
Harvard University	38	University of Florida	37, 49, 69
Illinois Institute of Technology	31, 44	University of Hawaii Manoa	37
Indiana University	5	University of Illinois	5, 23, 38, 44, 58
International Computer Science Institute	61	University of Iowa	13
Iowa State University	6, 25	University of Kansas	16
Johns Hopkins University	25, 32, 45	University of Maryland	52, 69, 70
Lehigh University	50	University of Massachusetts - Amherst	7
Loyola College in Maryland	74	University of Michigan	13, 25, 62
Massachusetts Institute of Technology	6, 25, 45, 57	University of Minnesota	13, 14, 26, 53
.....	58, 59, 62, 72	University of North Carolina	50, 63
Michigan State University	13, 72	University of North Carolina - Charlotte	8, 71
Michigan Technological University	46, 70	University of Notre Dame	6, 24, 45
New Jersey Institute of Technology	32	University of Pennsylvania	50, 73
New York University	59, 60	University of Pittsburgh	9, 28
North Carolina State University	71	University of Rhode Island	28
Northeastern University	32	University of Rochester	8
Northwestern University	5	University of South Florida	31
Ohio State University	27, 41	University of Southern California	4, 12, 22, 39, 43, 57
Oregon State University	27, 33, 38	University of Texas at Austin	14, 51
Princeton University	7, 14, 26, 59, 62	University of Utah	10, 29, 41, 58, 64, 73
Purdue University	6, 24, 49	University of Vermont	29
Rensselaer Polytechnic Institute	8	University of Virginia	10, 52, 58
Rutgers University New Brunswick	32	University of Washington	10, 11, 34, 47, 60
Salk Institute for Biological Studies	65	University of Wisconsin	14, 34, 48, 60, 61, 71
San Jose State University	30	Virginia Polytechnic Institute and State University	10, 30
Southern Illinois University at Carbondale	12	Washington State University	39, 54
Southern Methodist University	9	West Virginia University	73
Stanford University	3, 11, 42, 48, 52, 58, 72	William Marsh Rice University	41, 54
State University of New York - Binghamton	26, 70		
State University of New York - Stony Brook	8, 16, 49, 53		